arXiv:1611.10333v1 [math.ST] 30 Nov 2016

PCA from noisy, linearly reduced data: the diagonal case

Edgar Dobriban*, William Leeb†, and Amit Singer‡

Abstract

Suppose we observe data of the form $Y_i = D_i(S_i + \varepsilon_i) \in \mathbb{R}^p$ or $Y_i = D_i S_i + \varepsilon_i \in \mathbb{R}^p$, $i = 1, \dots, n$, where $D_i \in \mathbb{R}^{p \times p}$ are known diagonal matrices, ε_i are noise, and we wish to perform principal component analysis (PCA) on the unobserved signals $S_i \in \mathbb{R}^p$. The first model arises in missing data problems, where the D_i are binary. The second model captures noisy deconvolution problems, where the D_i are the Fourier transforms of the convolution kernels. It is often reasonable to assume the S_i lie on an unknown low-dimensional linear space; however, because many coordinates can be suppressed by the D_i, this low-dimensional structure can be obscured.

We introduce *diagonally reduced spiked covariance models* to capture this setting. We characterize the behavior of the singular vectors and singular values of the data matrix under high-dimensional asymptotics where $n, p \to \infty$ such that $p/n \to \gamma > 0$. Our results have the most general assumptions to date *even without diagonal reduction*. Using them, we develop optimal eigenvalue shrinkage methods for covariance matrix estimation and optimal singular value shrinkage methods for data denoising.

Finally, we characterize the error rates of the empirical Best Linear Predictor (EBLP) denoisers. We show that, perhaps surprisingly, their optimal tuning depends on whether we denoise in-sample or out-of-sample, but the optimally tuned mean squared error is the same in the two cases.

1 Introduction

Principal component analysis (PCA) is a classical statistical method that decomposes a collection of datapoints $s_1, \dots, s_n \in \mathbb{R}^p$ as a linear combination of vectors that account for the most variability (e.g., Jolliffe, 2002; Anderson, 2003). More formally, if $s_1, \dots, s_n$ are drawn from a probability distribution with mean zero and covariance matrix Σ_S, then the principal components (PCs) of the distribution are the eigenvectors $u_1, \dots, u_p$ of Σ_S. Typically, we approximate the distribution by projecting it onto the PCs with the largest eigenvalues.

*Department of Statistics, Stanford University. E-mail: dobriban@stanford.edu. Supported in part by NSF grant DMS-1407813, and by an HHMI International Student Research Fellowship.

†Program in Applied and Computational Mathematics, Princeton University. E-mail: wleeb@math.princeton.edu. Supported by the Simons Collaborations on Algorithms and Geometry.

‡Department of Mathematics, and Program in Applied and Computational Mathematics, Princeton University. E-mail: amits@math.princeton.edu. Partially supported by Award Number R01GM090200 from the NIGMS, FA9550-12-1-0317 from AFOSR, Simons Foundation Investigator Award and Simons Collaborations on Algorithms and Geometry, and the Moore Foundation Data-Driven Discovery Investigator Award.

A more specific model arising in many applications is the *spiked covariance model* (Johnstone, 2001). First, the signal S_i is a linear combination of r fixed but unobserved orthonormal PCs u_k:

$$S_i = \sum_{k=1}^{r} \ell_k^{1/2} z_{ik} u_k, \tag{1}$$

where r is a fixed parameter (independent of n and p) and z_{ik} are iid standardized random variables. Here ℓ_k are the eigenvalues, or equivalently the variances along the PCs u_k. Second, the observations are $X_i = S_i + \varepsilon_i$, where ε_i is noise with iid standardized entries. The spiked covariance model has been widely studied in probability and statistics (e.g., Baik et al., 2005; Baik and Silverstein, 2006; Paul, 2007; Benaych-Georges and Nadakuditi, 2012, etc); see also Paul and Aue (2014); Yao et al. (2015).

This paper considers the setting when the vectors S_i are not only corrupted by noise, but also *linearly reduced*. This means that for given matrices $D_i \in \mathbb{R}^{q_i \times p}$, we observe either

$$Y_i = D_i X_i = D_i S_i + D_i \varepsilon_i. \tag{2}$$

or

$$Y_i = D_i S_i + \varepsilon_i. \tag{3}$$

We think of the *reduction matrix* D_i as either a projection matrix or a linear filter reducing the information that we observe. In general, it will not be possible to reconstruct a vector v from $D_i v$.

We refer to model (2) as the *reduced-noise model*, and to model (3) as the *unreduced-noise model*. In the reduced-noise model, both the signal and noise are reduced, while in the unreduced-noise case, only the signal is. These models generalize the spiked covariance model, and arise naturally in several settings. For instance:

1. **Missing data:** For diagonal matrices D_i with zeros or ones the reduced-noise model from (2) corresponds to missing data problems widely encountered in statistics (e.g., Schafer, 1997; Little and Rubin, 2014). For random D_i independent of other variables, we are under the assumption of missing completely at random (MCAR).

2. **Deconvolution and image restoration:** In image processing, an image might be corrupted by "blurring"—convolution with a linear filter—followed by noise. After taking the Fourier transform, this can be modeled as a coordinate-wise multiplication by a diagonal matrix D_i, followed by adding noise. This corresponds to the unreduced-noise model from (3). For example, the image formation model in cryo-electron microscopy (cryo-EM) under the linear, weak phase approximation leads to such a model (Frank, 1996). A closely related model was recently used by Bhamre et al. (2016), where S_i are Fourier transforms of projection images of molecules, and D_i are contrast transfer functions.

3. **Structural variability in cryo-EM:** In cryo-EM, S_i is the three-dimensional structure of a molecule, and D_i is a tomographic projection of this volume onto a randomly selected plane. Since the molecule typically has only a few degrees of freedom, such as different conformations or states, it is reasonable to model S_i to lie on some unknown low-dimensional space (e.g., Katsevich et al., 2015; Andén et al., 2015). This corresponds to the unreduced-noise model from (3): the data $Y_i = D_i S_i + \varepsilon_i$ are tomographic projections $D_i S_i$ with added noise ε_i.

4. **Signal acquisition and compressed sensing:** In some signal acquisition tasks such as hyperspectral imaging, due to resource constraints it is convenient to acquire reduced or compressed measurements of signals. The reductions are often taken to be random projections. It is of interest to reconstruct the PCs of the original measurements (e.g., Chang, 2003; Fowler, 2009). This falls under the reduced-noise model from (2).

There may certainly be many other applications fitting this framework. In the above examples it is natural to posit that the distribution of the signal vectors S_i is of low effective dimensionality. In this paper we will assume that the distribution lies on some unknown linear space of small dimension r, as in equation (1). Given observations of the form (2) or (3), we address several natural statistical questions:

1. **Covariance estimation:** How to estimate the covariance matrix of the signals S_i? This is both a fundamental statistical problem, and has numerous applications, including classification and denoising.

2. **PCA:** How should we estimate the principal components of S_i? This question is of special interest due to the importance of PCA for exploratory data analysis and visualization.

3. **Denoising:** How can we denoise—or predict—the individual signal vectors S_i? This is a central question both in the missing data problems, where it corresponds to imputation, as well as in the image processing problems, where it amounts to noise reduction.

In this paper, we develop new methods for a special class of models, where the matrices $D_i \in \mathbb{R}^{p\times p}$ are *diagonal*. We will call the observations Y_i from (2) or (3) *diagonally reduced*. The missing data and deconvolution problems belong to this class. In the high-dimensional asymptotic regime where p, n both grow to infinity and $p/n \to \gamma > 0$, we develop methods that provide clear, quantitative answers to all questions posed above, under quite weak assumptions.

Related work by Katsevich et al. (2015) and Andén et al. (2015) develops methods for covariance estimation when the D_i's are projection matrices mapping a 3-D electron density to its integral on a randomly chosen plane. In this data acquisition model for cryo-EM, the authors propose consistent estimators of the covariance of the electron density. However, their observation models are different from our diagonally reduced models. In Bhamre et al. (2016), the questions of covariance estimation and denoising are studied empirically when the D_i's come from the contrast transfer function of a microscope and the S_i's are Fourier transforms of clean tomographic projections.

Nadakuditi (2014) develops methods for low-rank matrix estimation with missing data. Our results are more general, and also include methods for *covariance estimation* and *denoising individual datapoints*, see Sec. 2.3 for more details. Lounici (2014) develops eigenvalue soft thresholding methods for covariance estimation with missing data. In our somewhat more specialized models, we instead find the *optimal eigenvalue shrinkers*, and they are different from soft thresholding. Cai and Zhang (2016) develop minimax rate-optimal covariance matrix estimators for missing data, focusing on bandable and sparse models. We instead focus on the spiked covariance model.

We next give a brief overview of our results.

1.1 Probabilistic Results

A lot of work in random matrix theory studies the asymptotic spectral theory of the spiked covariance model and its variants, see e.g., Paul and Aue (2014); Yao et al. (2015). In Sec. 2, we

introduce two general *diagonally reduced* spiked covariance models corresponding to (2) and (3). We characterize the limiting eigenvalues of the data matrix Y (with rows $Y_i^\top$), and the limiting angles of its singular vectors with the population singular vectors (of the matrix S with rows $S_i^\top$).

More specifically, we show in Thm. 2.1 that the eigenvalue distribution of $n^{-1}Y^\top Y$ converges to a general Marchenko-Pastur distribution (Marchenko and Pastur, 1967), while the top few eigenvalues have well-defined almost sure limits. This mirrors the behavior known in unreduced spiked models, (e.g., Baik et al., 2005; Baik and Silverstein, 2006; Benaych-Georges and Nadakuditi, 2012); however, our assumptions in Sec. 2 are very general, and in fact lead to the *most general results to date even in the unreduced case* when $D_i = I_p$ (see Cor. 2.3 and Sec. 2.3 for discussion).

In the special case where the entries of D_i are iid, the limiting spectrum and the angles between the population and empirical singular vectors are given by explicit formulas related to the standard Marchenko-Pastur law, as described in Cor. 2.2. For general D_i, we can compute numerically the quantities specified by Thm. 2.1 with the SPECTRODE method (Dobriban, 2015); see Sec. 2.4. All computational results of this paper are reproducible with software publicly available at `github.com/dobriban/diagonally_reduced/`.

1.2 Covariance estimation

In Sec. 3, we apply the probabilistic results from Sec. 2 to develop methods for covariance estimation from diagonally reduced data, under the additional assumption that the diagonal entries of the reduction matrices D_i are iid. The prototypical example is the missing data problem with uniform missingness, where in the reduced-noise model (2) the entries are Bernoulli(δ).

Already for unreduced data from the spiked covariance model, the sample covariance matrix is a poor estimator of the population covariance, since neither the empirical PCs nor the empirical spectrum converge to their population counterparts. Though little can be done about correcting the PCs, one can develop optimal shrinkage estimators of the eigenvalues. In the unreduced case, Donoho et al. (2013) consider estimators of the form $U\eta(\Lambda)U^\top$, where U is the orthogonal matrix of PCs, Λ is the matrix of eigenvalues of the sample covariance, $\eta : \mathbb{R} \to \mathbb{R}$ is a shrinkage function, and $\eta(\Lambda)$ replaces every diagonal element λ of Λ by $\eta(\lambda)$.

In Sec. 3, we take a similar approach to the problem of covariance estimation in reduced models. We first define an unbiased estimator of the population covariance of the unreduced signals (Eq. 10). Building on the results of Sec. 2, we describe the asymptotic spectral theory of this estimator, and finally derive shrinkers of the spectrum that are asymptotically optimal for certain loss functions. We explicitly derive the optimal shrinkers in the case of operator norm loss (Eq. (15) for reduced-noise and Eq. (23) for unreduced-noise) and Frobenius norm loss (Eq. (18) for reduced-noise and Eq. (24) for unreduced-noise). We also derive the asymptotic errors of the optimal shrinkers (Eq. (17) and (19)), and give a recipe for deriving the optimal shrinkers for a broad class of loss functions.

1.3 Denoising

In Sec. 4, we consider denoising, the task of predicting the signal vectors S_i based on the observations Y_i. For this we study the Best Linear Predictor (BLP) well-known from random effects models (e.g., Searle et al., 2009). The general form of a BLP is $\hat{S}_i = \Sigma_S(\Sigma_S + \Sigma_\varepsilon)^{-1}Y_i$, where Σ_ε is the covariance matrix of the noise. This is also known as a "linear Bayesian" method (Hartigan, 1969). In other areas such as electrical engineering and signal processing, it is known as the "Wiener

filter", "(linear) Minimum Mean Squared Estimator (MMSE)" (Kay, 1993, Ch. 12), "linear Wiener estimator" (Mallat, 2008, p. 538), or "optimal linear filter" (MacKay, 2003, p. 550-551).

The BLP is an "oracle" method, because it depends on unknown population parameters. In practice we can use the empirical BLP (EBLP), $\hat{\Sigma}_S(\hat{\Sigma}_S + \hat{\Sigma}_\varepsilon)^{-1}Y_i$, where the unknown parameters are estimated using the data. Due to the inconsistency of PCA in high dimensions, this is sub-optimal to the BLP. However, we can find the asymptotically optimal method for estimating the covariance matrix Σ_S using the empirical PCs. This estimator holds several surprises. In particular, the optimal EBLP coefficients are different for *in-sample* and *out-of-sample* denoising—but the optimal mean squared error ends up identical! See Thms. 4.1, 4.3 for reduced-noise and Sec. 4.4 for unreduced-noise. It also turns out that the formula for in-sample EBLP, applied to all Y_i, is identical to optimal singular value shrinkage estimators (Sec. 4.2.2).

This analysis involves characterizing random quantities with an intricate dependence structure, such $Y_i^\top D_i \hat{u}_k$, where $\hat{u}_k$ is the k-th PC of the sample covariance matrix of Y_i. For this we extend significantly the technique introduced by Benaych-Georges and Nadakuditi (2012) to study the angles between u_k and $\hat{u}_k$. We call this approach the *outlier equation* method (see Sec. 4).

2 Probabilistic results

2.1 Main probabilistic results

This section presents a new result in random matrix theory, which will be the key tool for our work on covariance estimation. Recall that we have diagonally reduced observations $Y_i = D_i(S_i + \varepsilon_i)$ or $Y_i = D_i S_i + \varepsilon_i,\ i = 1, \ldots, n$, where $S_i \in \mathbb{R}^p$ are unobserved signals and $D_i \in \mathbb{R}^{p \times p}$ are diagonal matrices. The signals have the form $S_i = \sum_{k=1}^r \ell_k^{1/2} z_{ik} u_k$.

Here u_k are deterministic signal directions with $\|u_k\| = 1$. We will assume that u_k are delocalized, so that $|u_k|_\infty \leq C \log(p)^B / p^{1/2}$ for some constants $B, C > 0$. The scalars z_{ik} are standardized independent random variables, specifying the variation in signal strength from sample to sample. For simplicity we assume that the deterministic spike strengths are different and sorted: $\ell_1 > \ell_2 > \ldots > \ell_r > 0$. Finally $\varepsilon_i = \Gamma^{1/2}\alpha_i$ is sampling noise, where $\Gamma = \mathrm{diag}(g_1^2, \ldots, g_p^2)$ is diagonal and deterministic, and $\alpha_i = (\alpha_{i1}, \ldots, \alpha_{ip})^\top$ has independent standardized entries.

The diagonal matrices $D_i = \mathrm{diag}(D_{i1}, \ldots, D_{ip})$ have the form

$$D_i = \mu + \Sigma^{1/2} E_i, \tag{4}$$

where E_i have independent standardized entries. Here $\mu = \mathrm{diag}(\mu_1, \ldots, \mu_p) \in \mathbb{R}^{p \times p}$ is the deterministic diagonal mean and $\Sigma = \mathrm{diag}(\sigma_1^2, \ldots, \sigma_p^2) \in \mathbb{R}^{p \times p}$ is the deterministic diagonal covariance matrix of the entries of the reduction matrices. Let H_p be the uniform distribution on the p scalars $g_j^2 \cdot \mathbb{E}D_{ij}^2 = g_j^2 \cdot (\mu_j^2 + \sigma_j^2)$, $j = 1, \ldots, p$, and let G_p be the analogous object for g_j^2, $j = 1, \ldots, p$. It turns out that these are the distributions of noise variances relevant for our models. For dealing with model (2), we assume that as $p \to \infty$, H_p converges to a compactly supported limit distribution H: $H_p \Rightarrow H$. For dealing with model (3), we assume that G_p converges to a compactly supported limit distribution G: $G_p \Rightarrow G$.

We will consider the high-dimensional regime where $n, p \to \infty$ such that $p/n \to \gamma > 0$. In this setup, our answers will depend on the general Marchenko-Pastur distribution (Marchenko and Pastur, 1967). This law describes the behavior of empirical eigenvalue distributions of sample covariance matrices: If N is an $n \times p$ matrix with iid standardized entries, and T is a $p \times p$ positive

Table 1: Definitions

Name	Definition	Defined in
Reduced Observation	$Y_i = D_i(S_i + \varepsilon_i)$ or $Y_i = D_i S_i + \varepsilon_i$	(2), (3)
Signal	$S_i = \sum_{k=1}^{r} \ell_k^{1/2} z_{ik} u_k$	(1)
Reduction Matrix	$D_i = \mu + \Sigma^{1/2} E_i$	(4)
	$\mu = \mathrm{diag}(\mu_1, \ldots, \mu_p), \quad \Sigma = \mathrm{diag}(\sigma_1^2, \ldots, \sigma_p^2)$	
Noise	$\varepsilon_i = \Gamma^{1/2} \alpha_i$, where $\Gamma = \mathrm{diag}(g_1^2, \ldots, g_p^2)$	
Noise Variances	$H_p = p^{-1} \sum_{j=1}^{p} \delta_{g_j^2(\mu_j^2 + \sigma_j^2)}$, $H_p \Rightarrow H$	
	$G_p = p^{-1} \sum_{j=1}^{p} \delta_{g_j^2}$, $G_p \Rightarrow G$	
General MP Law	$F_{\gamma,H}, \quad \underline{F}_{\gamma,H}(x) = \gamma F_{\gamma,H}(x) + (1-\gamma)\delta_0$	
Stieltjes Transform	$m_{\gamma,H}(z) = \int \frac{dF_{\gamma,H}(x)}{x-z}, \quad \underline{m}_{\gamma,H}(z) = \int \frac{d\underline{F}_{\gamma,H}(x)}{x-z}$	
	$m_H(z) = \int \frac{dH(x)}{x-z}$	
D-Transform	$D_{\gamma,H}(x) = x \cdot m_{\gamma,H}(x) \cdot \underline{m}_{\gamma,H}(x)$	
Upper Edge	$b_H^2 = \sup \mathrm{supp}(F_{\gamma,H})$	

semidefinite matrix with eigenvalue distribution converging to H, then the eigenvalue distribution of the $p \times p$ matrix $n^{-1}T^{1/2}N^\top N T^{1/2}$ converges almost surely (a.s.) to the Marchenko-Pastur distribution $F_{\gamma,H}$ (see e.g., Bai and Silverstein, 2009, for a reference).

When $T = I_p$ is the identity, $F_{\gamma,H}$ is known as the *standard Marchenko-Pastur distribution*, and has density (if $\gamma \in (0,1)$):

$$f_\gamma(x) = \frac{\sqrt{(g_+ - x)(x - g_-)}}{2\pi x} I(x \in [g_-, g_+])$$

where $g_\pm = (1 \pm \sqrt{\gamma})^2$. For general H, $F_{\gamma,H}$ does not have a closed form, but it can be studied numerically (see e.g., Dobriban, 2015).

Closely related to $F_{\gamma,H}$ is the distribution $\underline{F}_{\gamma,H}(x) = \gamma F_{\gamma,H}(x) + (1-\gamma)\delta_0$. This is the limit of the eigenvalue distribution of the $n \times n$ matrix $n^{-1}N^\top T N$. We will also need the Stieltjes transform $m_{\gamma,H}$ of $F_{\gamma,H}$, $m_{\gamma,H}(z) = \int (x-z)^{-1} dF_{\gamma,H}(x)$, and the Stieltjes transform $\underline{m}_{\gamma,H}$ of $\underline{F}_{\gamma,H}$. Based on these, one can define the D-transform of $F_{\gamma,H}$ by

$$D_{\gamma,H}(x) = x \cdot m_{\gamma,H}(x) \cdot \underline{m}_{\gamma,H}(x).$$

Up to the change of variables $x = y^2$, this agrees with the D-transform defined in Benaych-Georges and Nadakuditi (2012). Let b_H^2 be the supremum of the support of $F_{\gamma,H}$, and $D_{\gamma,H}(b_H^2) = \lim_{t \downarrow b} D_{\gamma,H}(t^2)$. It is easy to see that this limit is well defined, and is either finite or $+\infty$. Let us denote the support of a distribution H on $\mathbb{R}$ by $\mathrm{supp}(H)$.

Denote the normalized data matrix $\tilde{Y} = n^{-1/2}Y$, with the $n \times p$ matrix Y having rows $Y_i^\top$. Our main probability result, proved later in Sec. 5.1, is the following.

Theorem 2.1 (Diagonally reduced spiked models)**.** *Consider the observation models* (2) *and* (3)*, under the above assumptions. Suppose that*

1. $\mathbb{E}\alpha_{ij}^4 < C$, $\mathbb{E}E_{ij}^4 < C$, *and* $\mathbb{E}|z_i|^{4+\phi} < C$ *for some* $\phi > 0$ *and* $C < \infty$.
2. *Under model* (2), $\sup \operatorname{supp}(H_p) \to \sup \operatorname{supp}(H)$. *Under model* (3), $\sup \operatorname{supp}(G_p) \to \sup \operatorname{supp}(G)$.
3. *The squared norms* $\|\mu u_k\|^2$ *converge to* $\tau_k > 0$.
4. *Under model* (2), μu_k *are* generic *with respect to* $M = \Gamma(\Sigma + \mu^2)$ *in the sense that*

$$u_j^\top \mu (M - zI_p)^{-1} \mu u_k \to I(j = k) \cdot \tau_k \cdot m_H(z)$$

for all $z \in \mathbb{C}^+$, *where* m_H *is the Stieltjes transform of* H.

Under model (3), μu_k *are generic with respect to* $M = \Gamma$, *i.e.,* $u_j^\top \mu (M - zI_p)^{-1} \mu u_k \to I(j = k) \cdot \tau_k \cdot m_G(z)$.

Then

1. *Under model* (2), *the eigenvalue distribution of* $\tilde{Y}^\top \tilde{Y}$ *converges to the general Marchenko-Pastur law* $F_{\gamma,H}$ *a.s. In addition, the* k*-th largest singular value of* $\tilde{Y}$ *converges,* $\sigma_k(\tilde{Y}) \to t_k > 0$ *a.s., where*

$$t_k^2 = \begin{cases} D_{\gamma,H}^{-1}(\frac{1}{\tau_k \ell_k}) & \text{if} \quad \ell_k > 1/[\tau_k D_{\gamma,H}(b_H^2)], \\ b_H^2 & \text{otherwise.} \end{cases} \tag{5}$$

Moreover, let $\nu_j = \mu u_j / \|\mu u_j\|$ *be the* normalized reduced signals *and let* $\hat{u}_k$ *be the right singular vector of* $\tilde{Y}$ *corresponding to* $\sigma_k(\tilde{Y})$. *Then* $(\nu_j^\top \hat{u}_k)^2 \to c_{jk}^2$ *a.s., where*

$$c_{jk}^2 = \begin{cases} \frac{m_{\gamma,H}(t_k^2)}{D'_{\gamma,H}(t_k^2)\tau_k \ell_k} & \text{if} \quad j = k \quad \text{and} \quad \ell_k > 1/[\tau_k D_{\gamma,H}(b_H^2)], \\ 0 & \text{otherwise.} \end{cases}$$

2. *Under model* (3), *the analogous results hold with* G *replacing* H *everywhere.*

Assumption 4 needs explanation. This assumption generalizes the existing conditions for spiked models. In particular, it is easy to see that it holds when the vectors u_k are random with independent coordinates. Specifically, let x, y are two independent random vectors with iid zero-mean entries with variance $1/p$. Then $\mathbb{E}x^\top \mu (M - zI_p)^{-1} \mu x = p^{-1} \operatorname{tr} \mu (M - zI_p)^{-1} \mu$. Assumption 4 requires that this converges to $\tau \cdot m_H(z)$, which happens for instance when the vector μ itself has random independent coordinates with variance τ/p, or when it equals a multiple of the identity. However, Assumption 4 is more general, as it does not require any kind of randomness in u_k.

Thm 2.1 gives the limiting angles of the empirical eigenvectors $\hat{u}_k$ with the *reduced* population eigenvectors $\nu_j = \mu u_j / \|\mu u_j\|$. These are in general different from the true eigenvectors u_j. However, in our main application (Cor. 2.2 and the following sections) they are the same, because μ is a multiple of identity.

One can gain some insight into the result in the simpler case where the noise is uncolored, so that $\Gamma = I_p$. In that case, before reduction we have a spike strength ℓ and an average noise level of unity. After reduction under model (2), we have a spike strength $\ell \|\mu u\|^2$, and an average noise level $p^{-1} \sum_j \mathbb{E}D_{ij}^2 = p^{-1}(\|\mu\|^2 + \|\sigma\|^2)$, where $\sigma = (\sigma_1, \ldots, \sigma_p)$. For a delocalized u, we expect $\|\mu u\|^2 \approx p^{-1}\|\mu\|^2$. Therefore, *reduction in model* (2) *typically decreases the signal strength* by a factor of

$$\frac{\|\mu\|^2}{\|\mu\|^2 + \|\sigma\|^2}.$$

However, the spike strength after reduction—$\ell\|\mu u\|^2$—depends on the correlation between μ and u. In particular, the reduction can vary among the different PCs. In contrast it is easy to see in a similar way that reduction in model (3) may increase or decrease the signal strength.

The key strength of Thm. 2.1 is its generality. Specifically, there is essentially only one previous result on reduced spiked models, appearing in Nadakuditi (2014). However, that only studies iid Bernoulli projections under restrictive conditions on the noise, whereas we allow for (1) a general diagonal covariance structure Σ in the reduction matrices, as well as (2) a general diagonal noise structure Γ, and (3) more general moment conditions. Moreover, even in unreduced spiked models, our results are already the most general results to date (see Sec. 2.3).

We think that the generality is important for several reasons: first, for practical reasons it is good to have results that require as few assumptions as possible, especially unverifiable conditions like "randomness" in u_k or "orthogonal invariance" of the noise. As a consequence of these general results, existing tools like singular value shrinkage are shown to apply more generally, so this is a direct improvement. Second, from a theoretical perspective it is good to understand the reason for the "spiking" behavior; our results clarify for instance that "orthogonal invariance" of the noise is not needed.

2.1.1 Comments on the proof

The broad outline of the proof is inspired by the argument presented in Benaych-Georges and Nadakuditi (2012) for unreduced spiked models. However, there are several new steps. First, the proof in Benaych-Georges and Nadakuditi (2012) concerns only the unreduced case, and the dependence introduced by the random reduction matrices D_i is a new challenge. The observations in model (2) are $D_i X_i = D_i S_i + D_i \varepsilon_i$, so the "signal" $D_i S_i$ and "noise" $D_i \varepsilon_i$ are dependent. However, we show that the dependence is asymptotically negligible. For non-diagonal reduction matrices D_i, the depencence may be asymptotically non-negligible; this explains why we currently need the diagonal assumption.

As a second novelty, the proof involves finding the limits of certain quadratic forms $u_k^\top R(z) u_j$, where $R(z)$ is a specific resolvent matrix with complex argument z. Since the u_k are deterministic, the concentration arguments of Benaych-Georges and Nadakuditi (2012) are not available. Instead, we adapt the "deterministic equivalents" approach of Bai et al. (2007). For this we need to take the imaginary part of the complex argument z to zero, which appears to be a new argument in this context.

2.2 Reduced standard spiked models

We will later use the following corollary for the *reduced standard spiked model* where the reduction coefficients and the noise entries are iid random variables. Suppose that the noise variances are equal to unity, so $\Gamma = I_p$ and thus ε_i have independent standardized entries. Moreover assume that the reduction matrices D_i have iid random diagonal entries D_{ij} with mean $\mu = \mathbb{E}D_{ij}$ and variance $\sigma^2 = \mathrm{Var}[D_{ij}]$. Note that previously μ was a matrix, but from now on it will be a scalar, and there will be no possibility for confusion. Let $m = \mu^2 + \sigma^2$ and $\delta = \mu^2/m$. In this case it is easy to see that the reduced eigenvectors are the same as the unreduced ones, i.e., $\nu_k = u_k$. Moreover, Assumption 4 from Thm. 2.1 reduces to $u_k^\top u_l \to 0$ as $n \to \infty$, if $k \neq l$.

Our answers can be expressed in terms of the well-known characteristics of the standard spiked model. The asymptotic location of the top singular values will depend on *the spike forward map*

(e.g., Baik et al., 2005):

$$\lambda(\ell) = \lambda(\ell;\gamma) = \begin{cases} (1+\ell)\left(1+\frac{\gamma}{\ell}\right) & \text{if } \ell > \gamma^{1/2}, \\ (1+\gamma^{1/2})^2 & \text{otherwise.} \end{cases}$$

The asymptotic angle between singular vectors will depend on *the cosine forward map* $c(\ell;\gamma) \geq 0$ given by (e.g., Paul, 2007; Benaych-Georges and Nadakuditi, 2011, etc):

$$c(\ell)^2 = c(\ell;\gamma)^2 = \begin{cases} \frac{1-\gamma/\ell^2}{1+\gamma/\ell} & \text{if } \ell > \gamma^{1/2}, \\ 0 & \text{otherwise.} \end{cases} \tag{6}$$

Corollary 2.2 (Reduced standard spiked models)**.** *Under observation model* (2)*, in the above setting, the eigenvalue distribution of* $m^{-1}\tilde{Y}^\top\tilde{Y}$ *converges to the standard Marchenko-Pastur law with aspect ratio* γ*, a.s. Moreover,* $m^{-1/2}\sigma_k(\tilde{Y}) \to t_k > 0$ *a.s., where*

$$t_k^2 = \lambda(\delta\ell_k) \tag{7}$$

Finally, let $\hat{u}_k$ *be the right singular vector of* $\tilde{Y}$ *corresponding to* $\sigma_k(\tilde{Y})$*. Then* $(u_j^\top\hat{u}_k)^2 \to c_{jk}^2$ *a.s., where*

$$c_{jk}^2 = \begin{cases} c^2(\delta\ell_k) & \text{if } j = k \\ 0 & \text{otherwise.} \end{cases} \tag{8}$$

Under observation model (3) *in the above setting, we have* $\sigma_k(\tilde{Y})^2 \to t_k^2 = \lambda(\mu^2\ell)$*, while* $(u_j^\top\hat{u}_k)^2 \to c_{jk}^2$*, where* $c_{jk}^2 = c^2(\mu^2\ell_k)$ *if* $j = k$ *and 0 otherwise.*

For the proof, see Sec. 5.3. While in the current paper we only use this corollary of Thm. 2.1, the proof in the special case is essentially as involved as in the general case. For this reason, and for potential future applications, we prefer to state Thm. 2.1 as well.

In this special case, reduction in model (2) lowers the spike strength from ℓ to $\delta\ell$. This result is related to Thm 2.4 of Nadakuditi (2014) on missing data, but we have the following advantages: (1) our result works under a 4-th order moment condition instead of requiring all bounded moments; (2) our result admits arbitrary diagonal reductions, not just missing data (see Sec. 2.3 for details).

Finally, when there is no reduction, i.e., when $D_i = I_p$ for all i, then it turns out we do not need the delocalization of u_k. Indeed that is only needed to show that the diagonal reductions introduce a negligible amount of dependence, but we do not need this when $D_i = I_p$. Hence, we can state the following corollary for unreduced spiked models:

Corollary 2.3 (Standard spiked models)**.** *Suppose we observe unreduced signals* $Y_i = S_i + \varepsilon_i$*, and we do* not *assume the delocalization of the PCs* u_k*. Suppose that the other assumptions of Cor. 2.2 hold:* $S_i = \sum_{k=1}^r \ell_k^{1/2} z_{ik} u_k$*, where* $u_k^\top u_j \to \delta_{kj}$*, while* z_{ik}*,* ε_{ij} *are iid standardized with* $\mathbb{E}|z_i|^{4+\phi} < C$*,* $\mathbb{E}\varepsilon_{ij}^4 < C$ *for some* $\phi > 0$ *and* $C < \infty$*.*

Then the conclusions of Cor. 2.2 hold. Specifically, the spectrum of $n^{-1}Y^\top Y$ *converges to a standard MP law, its spikes converge a.s. to* $\lambda(\ell_k)$ *and the squared cosines between population and sample eigenvectors converge a.s. to* $c^2(\ell_k)$*.*

Again, the key strength of this corollary is its generality, specifically that it only has 4-th moment assumptions, not orthogonal invariance.

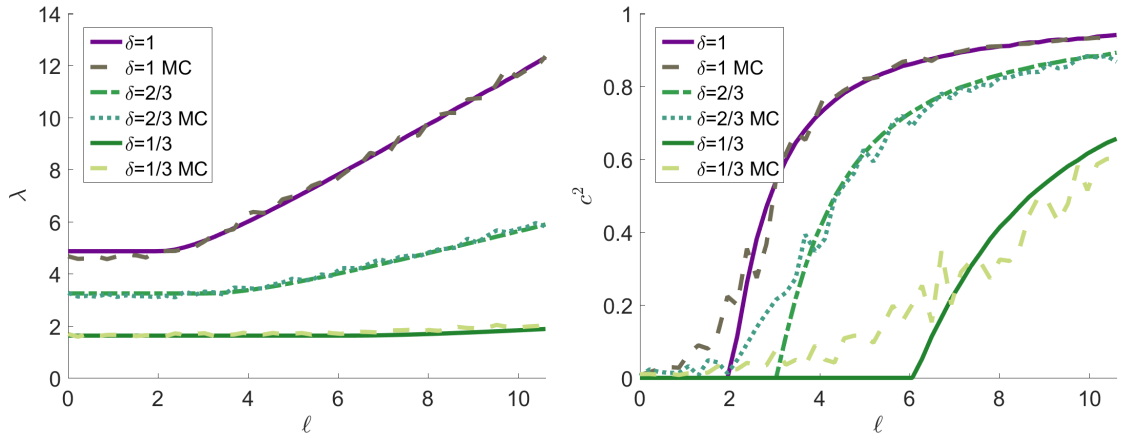

Figure 1: The effect of reduction on sample spikes and correlations between PCs. Formulas from Thm. 2.1 computed with SPECTRODE (Dobriban, 2015) overlaid with simulations. See Sec. 2.4.

2.3 Related work

There is substantial earlier work on unreduced spiked models, and even in this case our result leads to an improvement. We refer to Paul and Aue (2014); Yao et al. (2015) for general overviews of the area. The paper of Benaych-Georges and Nadakuditi (2012) is closely related to our approach, and we essentially follow their novel technique, relying on controlling certain bilinear forms. When the signal direction u is fixed, their results require the distribution of the noise matrix to be bi-unitarily invariant, which essentially reduces to Gaussian distributions. Our model is more general since it only requires a fourth-moment condition on the noise.

The technique introduced by Benaych-Georges and Nadakuditi (2011, 2012) was adapted to non-white Gaussian signal-plus noise matrices $X_i = \ell^{1/2} z_i u + \Gamma^{1/2} \varepsilon_i$ in Chapon et al. (2012). They rely on an integration by parts formula for functionals of Gaussian vectors and the Poincaré-Nash inequality. A Poincaré inequality was also assumed in Capitaine (2013). Our result is stronger, since we only require fourth moment conditions.

Previous extensions to the setting of missing data are found in Nadakuditi (2014), which describes the OptShrink method for singular value shrinkage matrix denoising. The method is extended to data missing at random, and in particular the limit spectrum of the data matrix with zeroed-out missing data is found (his Thm. 2.4). This is related to Thm. 2.1, but we have the following advantages: (1) our result works under the optimal 4-th order moment condition instead of requiring all moments to be bounded; (2) our result admits arbitrary reductions, not just binary projection matrices, (3) we extend to reduction matrices D_i that have non-iid entries, (4) we allow heteroskedastic diagonal noise $\varepsilon_i = \Gamma^{1/2} \alpha_i$; and (5) we also consider the unreduced-noise model from Eq. (3) (whereas Nadakuditi (2014) considers the reduced-noise model from Eq. (2)).

2.4 A numerical study

We report the results of a numerical study to gain insight into our theoretical results. We consider model (2), where $X_i = \ell^{1/2} z_i u + \varepsilon_i$, where the noise is heteroskedastic, $\varepsilon_i \sim \mathcal{N}(0, \Gamma)$, and Γ is

the diagonal matrix of eigenvalues of a $p \times p$ autoregressive (AR) covariance matrix of order 1, with entries $\Sigma_{ij} = \rho^{|i-j|}$. We set $Y_i = D_i X_i$, where D_i have iid Bernoulli(δ) entries. We set the missingness parameter δ to the values $1/3, 2/3$ and 1. We choose the AR autocorrelation coefficient $\rho = 0.5$, and vary the spikes ℓ from 0 to 3.5.

We compute numerically the formulas in Thm. 2.1, using the recent SPECTRODE method (Dobriban, 2015), see Sec. 5.1.6 for the details. We compare this with a Monte Carlo simulation with $n = 200$, $\gamma = 1/2$, z_i, u generated as Gaussian random variables, and the results averaged over 10 Monte Carlo trials. The results—displayed on Fig. 1—allow us to study the effect of reduction on spiked models. In particular, we observe the following phenomena:

- The theory and simulations show good agreement. For eigenvalues, the results are very accurate. For the cosine, the results are more variable, and especially so for small δ.
- In the left plot of Fig. 1, we see that the empirical spike is an increasing function of the population spike ℓ. Moreover, the location of the phase transition (PT) decreases with δ, i.e., reduction degrades the critical signal strength.
- Similarly, in the right plot of Fig. 1, we see that the cosine between population and empirical PCs increases with the population spike ℓ. For a given ℓ, the cosine decreases as $\delta \to 0$.

It is not hard, but beyond our scope, to formalize the last two observations into theorems.

3 Covariance matrix estimation

In this section, we develop methods for covariance estimation in the reduced-noise model $Y_i = D_i X_i + D_i \varepsilon_i$ (Secs. 3.2, 3.3) and in the unreduced-noise model $Y_i = D_i X_i + \varepsilon_i$ (Sec. 3.4). We also discuss some related work in Sec. 3.5. Finally, we present numerical experiments illustrating the results in Sec. 3.6.

We restrict our attention to a special case of the diagonally reduced model we considered in Sec. 2. We assume as in Sec. 2.2 that the entries of the reduction matrices D_i are independently and identically distributed. We also suppose that the noise ε_i is white, with variance 1 on each coordinate; that is, $\mathrm{Cov}(\varepsilon_i) = I_p$. We will also require that the coordinates of ε_i and the diagonal entries of D_i both have finite eighth moments. Recall that in the setting of Cor. 2.2, we have $\mu = \mathbb{E} D_{ij}$ and also that $\sigma^2 = \mathrm{Var}(D_{ij})$. The second moment is $m = \mathbb{E} D_{ij}^2 = \mu^2 + \sigma^2$, and $\delta = \mu^2/m$. For data missing uniformly at random, $m = \mu = \delta$ is the probability that each entry is observed.

3.1 The reduced-noise model

In the reduced-noise model, we observe n samples of the random vector $Y = D(S + \varepsilon)$. It is then easy to see that we have the following formulas relating the covariance matrix of the signal Σ_S and the covariance matrix of the observation Σ_Y:

$$\begin{aligned} \Sigma_Y &= \mu^2 \Sigma_S + \sigma^2 \operatorname{diag}(\Sigma_S) + m I_p \\ \Sigma_S &= \frac{1}{\mu^2} \Sigma_Y - \frac{\sigma^2}{m \mu^2} \operatorname{diag}(\Sigma_Y) - I_p \end{aligned} \tag{9}$$

These equations make it clear that the sample covariance matrix $\hat{\Sigma}_Y = n^{-1}\sum_{i=1}^n Y_i Y_i^\top$ is a biased estimator of the signal covariance matrix Σ_S. Based on the second equation, we consider the following *debiased* estimator of Σ_S:

$$\hat{\Sigma}_S = \frac{1}{\mu^2}\hat{\Sigma}_Y - \frac{\sigma^2}{m\mu^2}\operatorname{diag}(\hat{\Sigma}_Y) - I_p. \tag{10}$$

Here we assume for simplicity that μ, σ^2 are known; but these scalar parameters are straightforward to estimate from the observed D_i. In the special case of data missing completely at random, i.e., of iid sampling of entries with probability δ, we have $\mu^2 = \delta^2$ and $m = \delta$, so this formula becomes

$$\hat{\Sigma}_S = \frac{1}{\delta^2}\hat{\Sigma}_Y + \left(\frac{1}{\delta} - \frac{1}{\delta^2}\right)\operatorname{diag}(\hat{\Sigma}_Y) - I_p. \tag{11}$$

If our goal is to estimate $\Sigma_X = \Sigma_S + I_p$ instead of Σ_S, the corresponding unbiased estimator is $\hat{\Sigma}_X = \hat{\Sigma}_Y/\delta^2 + (\delta - 1)\operatorname{diag}(\hat{\Sigma}_Y)/\delta^2$. This recovers the unbiased estimator of Σ_X proposed by Lounici (2014). That paper proposes to estimate Σ_X by applying the soft-thresholding function $\eta_\tau(\lambda) = (\lambda - \tau)_+$ to the empirical eigenvalues λ of the covariance $\hat{\Sigma}_X$. Lounici (2014) proves error bounds for this estimator in both operator and Frobenius norm losses, for covariance matrices Σ_X of small effective rank $r_{eff}(\Sigma) = \operatorname{tr}(\Sigma)/\|\Sigma\|_{op}$. In contrast, we want to estimate the covariance matrix Σ_S of the signal. For this different task, in the spiked covariance model, the function η_τ is not optimal, as we will show in Section 3.3.

In the next section, we employ the probabilistic results from Sec. 2 to determine the asymptotic spectral theory of $\hat{\Sigma}_Y$. In particular, we find asymptotic formulas for the eigenvalues, and the angles between its PCs and those of the population covariance Σ_S. Next, we show how to use these results in conjunction with the theory of Donoho et al. (2013) to derive optimal non-linearities η of the spectrum of $\hat{\Sigma}_Y$ to estimate Σ_S for a variety of loss functions.

3.2 The asymptotic spectral theory of $\hat{\Sigma}_S$ in reduced-noise

In this section, we will analyze the asymptotic spectral theory of the debiased estimator $\hat{\Sigma}_S$; that is, the limiting eigenvalue distribution, spikes, and limiting angles of its top eigenvectors with those of Σ_S. We will rely on Corollary 2.2 from Section 2 and an argument controlling the diagonal terms in the proof in Sec. 5.4.1.

Corollary 3.1. *Let $1 \le k \le r$. Suppose that ℓ_k satisfies $\ell_k > \sqrt{\gamma}/\delta$. Then in the limit $p, n \to \infty$ and $p/n \to \gamma$, the k^{th} largest eigenvalue of $\hat{\Sigma}_S$, $k = 1, \dots, r$, converges almost surely to*

$$\frac{1}{\delta}(\delta\ell_k + 1)\left(1 + \frac{\gamma}{\delta\ell_k}\right) - \frac{1}{\delta}. \tag{12}$$

The distribution of the bottom $p - r$ eigenvalues of $\hat{\Sigma}_S$ converges to a shifted and scaled Marchenko-Pastur distribution $(\mu_{MP} - 1)/\delta$ supported on the interval $[(1 - \sqrt{\gamma})^2 - 1, (1 + \sqrt{\gamma})^2 - 1]/\delta$.

If $\hat{u}'_k$ is the k^{th} eigenvector of $\hat{\Sigma}_Y$ and $\hat{u}_k$ is the k^{th} eigenvector of $\hat{\Sigma}_S$, then almost surely we have $\lim_{n\to\infty}\langle \hat{u}'_k, \hat{u}_k\rangle^2 = 1$ and

$$\lim_{n\to\infty}\langle \hat{u}_k, u_k\rangle^2 = \frac{1 - \gamma/(\delta\ell_k)^2}{1 + \gamma/(\delta\ell_k)}.$$

Table 2: Optimal covariance shrinkage in unreduced-noise model. $\tilde{\ell}$ is the function defined by equation (16), $\delta = \mu^2/m$, and $s_k = \sqrt{1-c_k^2}$ is the asymptotic sine of the angle between the empirical and population PCs.

Loss function	Eigenvalue	Asymptotic loss	References
Operator	$\frac{\tilde{\ell}(\delta\lambda+1)}{\delta}$	$\ell_1 s_1$	(15), (17)
Squared Frobenius	$\frac{\tilde{\ell}(\delta\lambda+1)\cdot c^2(\tilde{\ell}(\delta\lambda+1))}{\delta}$	$\sum_{k=1}^{r}(1-c_k^4)\ell_k^2$	(18), (19)

If $\ell_k \leq \sqrt{\gamma}/\delta$, then the top eigenvalue converges to the upper edge of the shifted MP distribution, and the cosine converges to 0.

The shifted Marchenko-Pastur distribution arises as the limiting empirical spectral distribution of the eigenvalues corresponding to noise. This is also the case for the available case sample covariance of pure noise (Jurczak and Rohde, 2015). We will discuss the available-case estimator in Secs. 3.5 and 3.6.

3.3 Optimal shrinkage of the spectrum of $\hat{\Sigma}_S$ in reduced-noise

Having characterized the asymptotic spectrum of the debiased estimator $\hat{\Sigma}_S$, we can apply the technique of Donoho et al. (2013) to derive optimal shrinkers of the eigenvalues of $\hat{\Sigma}_S$ to minimize various loss functions. Any of the 26 loss functions found in Donoho et al. (2013) can be adapted to the setting of diagonally reduced data. In Sec. 5.4.2, we carefully check the details of this program.

Write the eigendecomposition of Σ_S as $\Sigma_S = U\Lambda U^\top$ and the eigendecomposition of the debiased estimator $\hat{\Sigma}_S$ as $\hat{\Sigma}_S = \hat{U}\hat{\Lambda}\hat{U}$. For a given function $\eta : \mathbb{R} \to [0,\infty)$, define the matrix $\hat{\Sigma}_S^\eta$ by

$$\hat{\Sigma}_S^\eta = \hat{U}\eta(\hat{\Lambda})\hat{U}^\top$$

where $\eta(\hat{\Lambda})$ is the diagonal matrix that replaces the k^{th} diagonal element $\hat{\lambda}_k$ of $\hat{\Lambda}$ with $\eta(\hat{\lambda}_k)$.

For any value of p, let $L_p(A,B)$ denote a loss function between two p-by-p symmetric matrices A and B. We consider loss functions with two key properties: first, they must be *orthogonally invariant*; that is, $L_p(A,B) = L_p(UAV,UBV)$ for any orthogonal matrices U and V. Second, they must *decompose over blocks*. This means that if $A_1, B_1 \in \mathbb{R}^{p_1\times p_1}$ and $A_2, B_2 \in \mathbb{R}^{p_2\times p_2}$ where $p = p_1 + p_2$, then either $L_p(A_1 \oplus A_2, B_1 \oplus B_2) = \max\{L_{p_1}(A_1,B_1), L_{p_2}(A_2,B_2)\}$, in which case we say L_p is *max-decomposable*; or $L_p(A_1 \oplus A_2, B_1 \oplus B_2) = L_{p_1}(A_1,B_1) + L_{p_2}(A_2,B_2)$, in which case we say L_p is *sum-decomposable*. Operator norm loss $L_p(A,B) = \|A-B\|_{op}$ is max-decomposable, whereas squared Frobenius norm loss $L_p(A,B) = \|A-B\|_F^2$ is sum-decomposable.

Our goal is to find the function η that minimizes the asymptotic loss over certain classes; that is, we seek:

$$\eta^* = \arg\min_\eta L_\infty(\Sigma_S, \hat{\Sigma}_S^\eta)$$

where $L_\infty(\Sigma_S, \hat{\Sigma}_S^\eta)$ is the almost sure limit of $L_p(\Sigma_S, \hat{\Sigma}_S^\eta)$ as $n,p \to \infty$ and $p/n \to \gamma$. We will show from first principles that this limit is well defined. As in Donoho et al. (2013), we will consider only those functions η that *collapse the vicinity of the bulk* to 0; that is, for which there is an $\varepsilon > 0$ such

that $\eta(\lambda) = 0$ whenever $\lambda \leq (1+\sqrt{\gamma})^2/\delta - 1/\delta + \varepsilon$ (this is the value of the upper bulk edge, as given in Cor. 3.1).

In Sec. 5.4.2, we will show:

$$L_\infty(\Sigma_S, \hat{\Sigma}_S^\eta) = L_{2r}\left(\bigoplus_{k=1}^{r} A_2(\ell_k), \bigoplus_{k=1}^{r} B_2(\eta(\lambda_k), c_k, s_k)\right) \tag{13}$$

where $A_2(\ell) = \begin{pmatrix} \ell & 0 \\ 0 & 0 \end{pmatrix}$, and $c_k = c(\delta\ell_k)$, $s_k = s(\delta\ell_k)$, where $s(\ell) = \sqrt{1-c^2(\ell)} \geq 0$ the asymptotic sine of the angle between the empirical and population PCs,

$$B_2(\eta(\lambda), c, s) = \begin{pmatrix} \eta(\lambda)c^2(\delta\ell) & \eta(\lambda)c(\delta\ell)s(\delta\ell) \\ \eta(\lambda)c(\delta\ell)s(\delta\ell) & \eta(\lambda)s^2(\delta\ell) \end{pmatrix}.$$

Since the loss function is either max-decomposable or sum-decomposable, for η to minimize the right side, it is sufficient that it minimize every individual term $L_2(A_2(\ell_k), B_2(\eta(\lambda_k), c_k, s_k))$. That is, the asymptotically optimal η minimizes the two-dimensional loss:

$$\eta^* = \operatorname*{argmin}_\eta L_2(A_2(\ell), B_2(\eta(\lambda), c, s)). \tag{14}$$

This dramatically simplifies the problem, as this minimization can often be done explicitly. Deriving the optimal η now depends on the particular choice of loss function. We consider two representative cases where a simple closed formula is easily found: operator norm loss, and squared Frobenius norm loss. The same recipe of explicitly solving the problem (14) can be used for any orthogonally-invariant and max- or sum-decomposable loss function, including those found in Donoho et al. (2013).

3.3.1 Operator norm loss/max-decomposable losses

Since operator norm loss $L_p(A, B) = \|A - B\|_{op}$ is max-decomposable, equation (13) implies that

$$L_\infty(\Sigma_S, \hat{\Sigma}_S^\eta) = \max_{1 \leq k \leq r} L_2(A(\ell_k), B(\eta(\lambda_k), c_k, s_k)).$$

Consequently, the asymptotically optimal η is the one that minimizes the two-dimensional loss function $L_2(A(\ell_k), B(\eta(\lambda_k), c_k, s_k)) = \|A(\ell_k) - B(\eta(\lambda_k), c_k, s_k)\|_{op}$. Repeating the derivation in Donoho et al. (2013), the optimal $\eta(\lambda_k)$ sends λ_k back to its population value, ℓ_k. From formula (12) in Cor. 3.1,

$$\eta^*(\lambda) = \frac{\tilde{\ell}(\delta\lambda + 1)}{\delta} \tag{15}$$

where $\tilde{\ell}$ inverts the spike forward map $\ell \mapsto \lambda(\ell; \gamma)$ defined in Sec. 2.2,

$$\tilde{\ell}(y) = \frac{y - 1 - \gamma + \sqrt{(y-1-\gamma)^2 - 4\gamma}}{2}. \tag{16}$$

Direct computation shows that $L_2(A(\ell_k), B(\eta^*(\lambda_k), c_k, s_k)) = \ell_k s_k$; consequently, the asymptotic loss is given by the formula:

$$L_\infty(\Sigma_S, \hat{\Sigma}_s^{\eta^*}) = \ell_1 s_1. \tag{17}$$

Table 3: Optimal covariance shrinkage in unreduced-noise model. $\tilde{\ell}$ is the function defined by equation (16), and $s_k = \sqrt{1-c_k^2}$ is the asymptotic sine of the angle between the empirical and population PCs.

Loss function	Eigenvalue	Asymptotic loss	References
Operator	$\frac{\tilde{\ell}(\mu^2\lambda+1)}{\mu^2}$	$\ell_1 s_1$	(23), (17)
Squared Frobenius	$\frac{\tilde{\ell}(\mu^2\lambda+1)\cdot c^2(\tilde{\ell}(\mu^2\lambda+1))}{\mu^2}$	$\sum_{k=1}^r (1-c_k^4)\ell_k^2$	(24), (19)

3.3.2 Frobenius norm loss/sum-decomposable losses

Since the squared Frobenius loss $L_p(A,B) = \|A-B\|_F^2$ is sum-decomposable, equation (13) implies that

$$L_\infty(\Sigma_S, \hat{\Sigma}_S^\eta) = \sum_{k=1}^r L_2(A(\ell_k), B(\eta(\lambda_k), c_k, s_k)).$$

Consequently, the asymptotically optimal η is the one that minimizes the two-dimensional loss function $L_2(A(\ell_k), B(\eta(\lambda_k), c_k, s_k)) = \|A(\ell_k) - B(\eta(\lambda_k), c_k, s_k)\|_F^2$. As derived in Donoho et al. (2013), the value of $\eta(\lambda_k)$ that minimizes this is $\ell_k c_k^2$. We have already seen that $\ell_k = \tilde{\ell}(\delta\lambda_k+1)/\delta$, where $\tilde{\ell}$ is the function defined by (16). Consequently, with $c^2(\ell) = c^2(\ell;\gamma)$ being the cosine forward map, the formula for $\eta^*(\lambda)$ is

$$\eta^*(\lambda) = \frac{\tilde{\ell}(\delta\lambda+1)\cdot c^2(\tilde{\ell}(\delta\lambda+1))}{\delta}. \tag{18}$$

A straightforward computation shows that $L_2(A(\ell_k), B(\ell_k c_k^2, c_k, s_k)) = (1-c_k^4)\ell_k^2$, and consequently, the asymptotic loss is given by the formula:

$$L_\infty(\Sigma_S, \hat{\Sigma}_s^{\eta^*}) = \sum_{k=1}^r (1-c_k^4)\ell_k^2. \tag{19}$$

3.4 The unreduced-noise model

In the unreduced-noise model $Y_i = D_i s_i + \varepsilon_i$ we can develop similar methods for covariance estimation. The formulas relating Σ_S to Σ_Y are

$$\begin{aligned} \Sigma_Y &= \mu^2\Sigma_S + \sigma^2 \operatorname{diag}(\Sigma_S) + I_p \\ \Sigma_S &= \frac{1}{\mu^2}\Sigma_Y - \frac{\sigma^2}{m\mu^2}\operatorname{diag}(\Sigma_Y) - \frac{1}{m}I_p. \end{aligned} \tag{20}$$

Consequently, the analogous *debiased* covariance estimator is:

$$\hat{\Sigma}_S = \frac{1}{\mu^2}\hat{\Sigma}_Y - \frac{\sigma^2}{m\mu^2}\operatorname{diag}(\hat{\Sigma}_Y) - \frac{1}{m}I_p. \tag{21}$$

In this section, we will derive optimal shrinkers of the spectrum of $\hat{\Sigma}_S$ in the unreduced-noise model using the same technique as for the reduced-noise model in Sec. 3.3. Our analysis rests on the following result, proved in Sec. 5.4.1:

Corollary 3.2. *Let $1 \le k \le r$. Suppose that ℓ_k satisfies $\ell_k > \sqrt{\gamma}/\mu^2$. Then in the limit $p, n \to \infty$ and $p/n \to \gamma$, the k^{th} largest eigenvalue of $\hat{\Sigma}_S$ converges almost surely to*

$$\frac{1}{\mu^2}(\mu^2\ell_k + 1)\left(1 + \frac{\gamma}{\mu^2\ell_k}\right) - \frac{1}{\mu^2}. \tag{22}$$

The distribution of the bottom $p - r$ eigenvalues of $\hat{\Sigma}_S$ converges to a shifted Marchenko-Pastur distribution supported on the interval $[(1-\sqrt{\gamma})^2 - 1, (1+\sqrt{\gamma})^2 - 1]/\mu^2$.

If $\hat{u}'_k$ is the k^{th} eigenvector of $\hat{\Sigma}_Y$ and $\hat{u}_k$ is the k^{th} eigenvector of $\hat{\Sigma}_S$, then almost surely we have $\lim_{n\to\infty}\langle \hat{u}'_k, \hat{u}_k\rangle^2 = 1$ and

$$\lim_{n\to\infty}\langle \hat{u}_k, u_k\rangle^2 = \frac{1 - \gamma/(\mu^2\ell_k)^2}{1 + \gamma/(\mu^2\ell_k)}.$$

If $\ell_k \le \sqrt{\gamma}/\mu^2$, then the top eigenvalue converges to the upper edge of the shifted MP distribution, and the cosine converges to 0.

Given an empirical eigenvalue $\hat{\lambda}_k$ of $\hat{\Sigma}_S$, we estimate the population eigenvalue ℓ_k by $\tilde{\ell}(\mu^2\hat{\lambda}_k + 1)/\mu^2$ where $\tilde{\ell}$ is the function given by formula (16). This estimator converges almost surely to the true value ℓ_k if ℓ_k exceeds the threshold $\sqrt{\gamma}/\mu^2$. This also gives us an estimator of the squared cosine, by the formula $c(\tilde{\ell}(\mu^2\hat{\lambda}_k + 1))$. We can now derive the optimal non-linear functions on the spectrum. For operator norm loss, we have

$$\eta^*(\lambda) = \frac{\tilde{\ell}(\mu^2\lambda + 1)}{\mu^2}, \tag{23}$$

which incurs an asymptotic loss of $L_\infty(\Sigma_S, \hat{\Sigma}_s^{\eta^*}) = \ell_1 s_1$. For squared Frobenius norm loss, the optimal non-linearity is

$$\eta^*(\lambda) = \frac{\tilde{\ell}(\mu^2\lambda + 1) \cdot c(\tilde{\ell}(\mu^2\lambda + 1))}{\mu^2} \tag{24}$$

and the asymptotic loss is $L_\infty(\Sigma_S, \hat{\Sigma}_s^{\eta^*}) = \sum_{k=1}^{r}(1 - c_k^4)\ell_k^2$.

3.5 Alternative linear systems for estimating Σ_S

The optimal shrinkers derived in Secs. 3.3 and 3.4 for estimating the covariance Σ_S from the reduced-noise observations start from the debiased estimators (10) for reduced-noise and (21) for unreduced noise. Another way of viewing these estimators is as the solution to a linear system: for reduced-noise, this system is given by equation (9), and for unreduced-noise by equation (20).

Of course, there are other linear systems yielding unbiased estimators whose spectrum we could shrink. The papers Katsevich et al. (2015); Andén et al. (2015); Bhamre et al. (2016) consider such an estimator, which we will briefly discuss here.

By definition, for any $k = 1, \ldots, n$, $\Sigma_S = \mathbb{E}_{S,\varepsilon}[S_k S_k^\top]$, where $\mathbb{E}_{S,\varepsilon}$ denotes the expectation with respect to the random S_k and ε_k, but not the D_k. We can therefore write in the unreduced-noise model $Y_k = D_k S_k + \varepsilon_k$:

$$D_k \Sigma D_k = \mathbb{E}_{S,\varepsilon}[(D_k S_k)(D_k S_k)^\top] = \mathbb{E}_{S,\varepsilon}[(Y_k - \varepsilon_k)(Y_k - \varepsilon_k)^\top] = \mathbb{E}_{S,\varepsilon}[Y_k Y_k^\top] + I_p. \tag{25}$$

If we knew the values of $\mathbb{E}_{S,\varepsilon}[Y_k Y_k^\top]$ for every k, we could derive an unbiased estimator of Σ_S by solving the n equations given by (25). The papers Katsevich et al. (2015); Andén et al. (2015); Bhamre et al. (2016) instead substitute the observed value $Y_k Y_k^\top$ for its expected value, and derive an unbiased estimator of Σ_S by the minimization problem

$$\hat{\Sigma}_S = \arg\min_{\Sigma} \frac{1}{n} \sum_{k=1}^{n} \|D_k \Sigma D_k^\top - Y_k Y_k^\top - I_p\|_F^2.$$

Differentiating in Σ, we see that $\hat{\Sigma}_S$ must satisfy the linear system:

$$\frac{1}{n} \sum_{k=1}^{n} D_k^\top D_k \Sigma D_k^\top D_k = \frac{1}{n} \sum_{k=1}^{n} (D_k^\top Y_k)(D_k^\top Y_k)^\top - \frac{1}{n} \sum_{k=1}^{n} D_k^\top D_k.$$

We now consider another linear system, defined by averaging the n equations (25):

$$\frac{1}{n} \sum_{k=1}^{n} D_k \Sigma_S D_k = \frac{1}{n} \sum_{k=1}^{n} \mathbb{E}_{S,\varepsilon}[Y_k Y_k^\top] - I_p.$$

Note that the matrices D_k are fixed in this equation. By replacing $\mathbb{E}_{S,\varepsilon}[Y_k Y_k^\top]$ with the estimate $Y_k Y_k^\top$, we define a new estimator of Σ_S as the solution to the equation

$$\frac{1}{n} \sum_{k=1}^{n} D_k \Sigma D_k = \frac{1}{n} \sum_{k=1}^{n} Y_k Y_k^\top - I_p. \tag{26}$$

In the case of reduced-noise, we can repeat the same derivation and arrive at an unbiased estimator that satisfies the system

$$\frac{1}{n} \sum_{k=1}^{n} D_k \Sigma D_k = \frac{1}{n} \sum_{k=1}^{n} Y_k Y_k^\top - \frac{1}{n} \sum_{k=1}^{n} D_k^2. \tag{27}$$

We will denote the estimator solving (26) (in the unreduced-noise case) and (27) (in the reduced noise case) by $\hat{\Sigma}'_s$. Taking the expectation of each side of (27), we arrive at the linear system (20), which defines our estimator $\hat{\Sigma}_S$ in the unreduced-noise model. Similarly, taking the expectation of each side of (26), we arrive at the linear system (9), which defines our estimator $\hat{\Sigma}_S$ in the reduced-noise model. Consequently, we expect that, in the limit $n, p \to \infty$, $\hat{\Sigma}_S$ and $\hat{\Sigma}'_S$ will be close. In fact, we can show that the relative error of the estimators converges to 0, as stated in the folowing proposition (proved in Sec. 5.4.3):

Proposition 3.3. *In both the reduced-noise and unreduced-noise models, the relative difference* $\|\hat{\Sigma}_S - \hat{\Sigma}'_S\|_F / \|\hat{\Sigma}_S\|_F \to 0$ *almost surely as* $n, p \to \infty$ *and* $p/n \to \gamma$.

3.6 Numerical experiments

We perform experiments with missing data, where the diagonal entries of the reduction matrices D_i are independent Bernoulli(δ) random variables. In this case, $\mu = m = \delta$, and $\sigma^2 = \delta(1-\delta)$. The unbiased estimator to which we apply shrinkage is given by the formula (11). In all experiments, both the signal and the noise are drawn from Gaussian distributions.

3.6.1 The errors in estimating the covariance

In the first experiment, we illustrate the dependence of the asymptotic errors on the parameters δ and γ. The clean signal vectors S_i are drawn from a rank 1 Gaussian. The noise is also Gaussian noise, of unit variance; the ambient dimension p is fixed at $p = 1200$ in this experiment, while the number of samples n varies with γ. The top two rows of Fig. 2 shows the errors in estimation for Frobenius loss and operator norm loss, as functions of the parameters γ and δ. The error bars cover the empirical mean error, plus/minus two standard deviations over 200 runs of the experiment.

Several phenomena are apparent in these plots. First, the empirical mean of the errors is well-approximated by the asymptotic error formulas (17) and (19), especially as the number of samples grows (corresponding to smaller γ, as $n = p/\gamma = 1200/\gamma$). Second, the errors decay as δ approaches 1, which is expected as larger δ increases the effective signal strength. Third, the errors grow as γ approaches 1; this is also expected, since large γ leads to higher dimensional problems.

4 Denoising

4.1 Setup

It is often of interest to denoise the observations Y_i and predict the signal components S_i. We envision a scenario where the data (Y_i, D_i), $i = 1, \ldots, n$ is already collected, and we construct the denoisers using this dataset. With *in-sample denoising* we denoise Y_i to predict the signal components S_i. This makes sense in many applications where we want to use the entire dataset to construct the denoiser.

A closely related scenario is *out-of-sample denoising*, where we want to denoise a new datapoint (Y_0, D_0). This arises in applications where new samples are made available after an initial pre-processing of $Y_1, \ldots, Y_n$ is performed, and it is not desired or not feasible to repeat this processing on the augmented dataset $(Y_0, D_0), (Y_1, D_1), \ldots, (Y_n, D_n)$ for every new data point.

While these two settings are very closely related, it turns out, perhaps surprisingly, that the optimal way to construct the denoisers differs substantially between the two. The reason turns out to be closely related to the observation that in high dimensions, the addition of a single datapoint changes the direction of the PCs. We will explain this phenomenon in detail below.

We first study the reduced-noise model from (2) in the setting of Cor. 2.2, where the diagonal entries of D_i are drawn iid from a distribution with mean μ and variance σ^2, and the observations are $Y_i = D_i(S_i + \varepsilon_i)$. This is a special type of *random effects* model, as the "effects" z_{ik} of the "factors" u_k in $S_i = \sum_{k=1}^{r} \ell_k^{1/2} z_{ik} u_k$ are random from sample to sample. As usual in random effects models, the optimal way to predict S_i from a mean squared error (MSE) perspective is to use the Best Linear Predictor—or BLP—(e.g., Searle et al., 2009, Sec. 7.4). The BLP of S_i is the predictor $\hat{S}_i = MY_i$ that minimizes $\mathbb{E}\|\hat{S}_i - S_i\|^2$.

It is well known that the BLP is $\hat{S}_i^{BLP} = \mathrm{Cov}\,[S_i, Y_i]\,\mathrm{Cov}\,[Y_i, Y_i]^{-1}\,Y_i$. Under the assumptions of Cor. 2.2, we can show (see Sec. 5.5) that the BLP has the same asymptotic MSE properties as

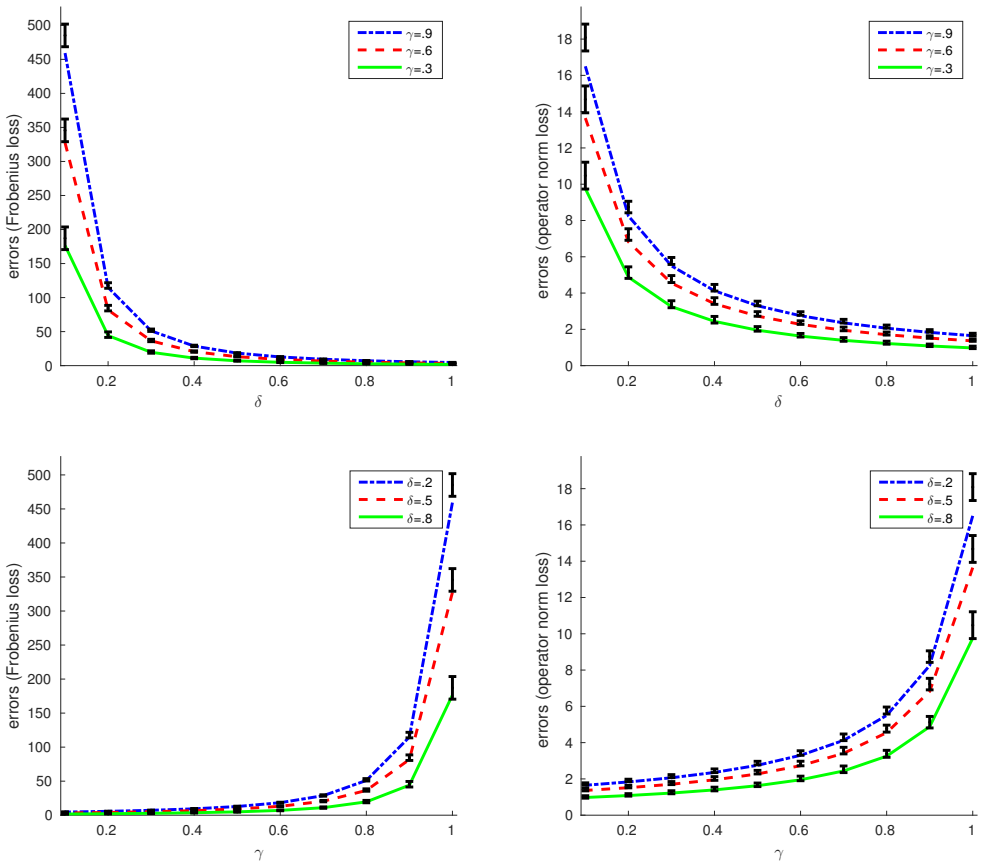

Figure 2: *Top row*: Estimation error of the covariance as a function of δ. Left: Frobenius shrinker. Right: operator norm shrinker. For each δ, 200 Monte Carlo tests were averaged. The error bars cover two standard deviations; the lines go through the predicted errors. *Bottom row*: Same plot as a function of γ.

Table 4: Denoising in the single-spiked case. BLP: population Best Linear Predictor using u. EBLP: in-sample empirical Best Linear Predictor using $\hat{u}$. EBLP-OOS: out-of-sample empirical Best Linear Predictor using $\hat{u}$. Here we abbreviate $\lambda = \lambda(\delta\ell;\gamma)$, $c^2 = c^2(\delta\ell;\gamma)$, $s^2 = 1-c^2$, $\beta = 1+\gamma/(\delta\ell)$.

Name	Definition	Asy MSE	Asy Opt η	Asy Opt MSE	Ref
BLP	$\eta \cdot uu^\top Y_i$	$(1-\eta\mu)^2\ell + \eta^2 m$	$\frac{\mu\ell}{\mu^2\ell+m}$	$\frac{m\ell}{\mu^2\ell+m}$	Thm. 4.1
EBLP	$\eta \cdot \hat{u}\hat{u}^\top Y_i$	$\ell + \eta^2 \cdot m\lambda - 2\eta \cdot \mu\ell c^2 \cdot \beta$	$\frac{\mu\ell c^2}{\mu^2\ell+m}$	$\ell \cdot \frac{\mu^2\ell c^2 s^2+m}{\mu^2\ell c^2+m}$	Thm. 4.1
EBLP-OOS	$\eta \cdot \hat{u}\hat{u}^\top Y_i$	$\ell + \eta^2 \cdot (\mu^2\ell c^2+m) - 2\eta \cdot \mu\ell c^2$	$\frac{\mu\ell c^2}{\mu^2\ell c^2+m}$	$\ell \cdot \frac{\mu^2\ell c^2 s^2+m}{\mu^2\ell c^2+m}$	Prop. 4.3

a denoiser of the following simpler form:

$$\hat{S}_i^{\tau,B} = \sum_{k=1}^{r} \tau_k u_k u_k^\top Y_i. \tag{28}$$

The denoisers $\hat{S}_i^{\tau,B}$ are indexed by $\tau = (\tau_1,\dots,\tau_r)$, and our argument shows that with the choice $\tau_k = \mu\ell_k/(\mu^2\ell_k+m)$ they are in fact asymptotically equivalent to BLP denoisers.

In practice, the true PCs u_k are not known, so we use the Empirical BLP (EBLP), where we estimate the unknown parameters u_k using the entire dataset. Here we will use the k-th top right singular vector $\hat{u}_k$ of the $n \times p$ matrix Y with rows $Y_i^\top$ as an estimator of u_k. In analogy with the simplified form of the denoisers in (28), we will consider EBLPs scaled by $\eta = (\eta_1,\dots,\eta_r)$ having the form:

$$\hat{S}_i^{\eta} = \sum_{k=1}^{r} \eta_k \hat{u}_k \hat{u}_k^\top Y_i. \tag{29}$$

Our goal is to find the optimal scalars η_k, and characterize their MSE.

4.2 In-sample denoising

First, we will study *in-sample* denoising, where the data to be denoised are also used to construct the denoisers. Our main findings for the asymptotic MSE (AMSE) are summarized in Table 4 (in the single-spiked case) and in the following theorem, proved in Sec. 5.6.

Theorem 4.1 (In-sample denoising)**.** *In the setting of Cor. 2.2 consider in-sample best linear predictors (BLP) of the signals S_i based on the observations Y_i.*

1. *The BLP denoisers $\hat{S}_i^{\tau,B} = \sum_{k=1}^{r} \tau_k \cdot u_k u_k^\top Y_i$ based on the population singular vectors have an AMSE $\lim_{n,p\to\infty} \mathbb{E}\|S_i - \hat{S}_i^{\tau,B}\|^2$ of*

$$AMSE^B(\tau_1,\dots,\tau_r;\ell_1,\dots,\ell_r,\gamma) = \sum_{k=1}^{r} AMSE^B(\tau_r;\ell_r,\gamma),$$

where $AMSE^B(\tau;\ell,\gamma) = (1-\tau\mu)^2\ell + \tau^2 m$ *is the AMSE of the BLP* $\hat{S}_i^{\tau,B} = \tau \cdot uu^\top Y_i$ *in a single-spiked model with spike strength* ℓ *under the assumptions of Cor. 2.2. The asymptotically optimal coefficients are*

$$\tau_k^* = \frac{\mu\ell_k}{\mu^2\ell_k + m}.$$

2. *The EBLP denoisers* $\hat{S}_i^\eta = \sum_{k=1}^r \eta_k \cdot \hat{u}_k\hat{u}_k^\top Y_i$*, based on the empirical singular vectors have an AMSE* $\lim_{n,p\to\infty} \mathbb{E}\|S_i - \hat{S}_i^\eta\|^2$ *of*

$$AMSE^E(\eta_1,\ldots,\eta_r;\ell_1,\ldots,\ell_r,\gamma) = \sum_{k=1}^r AMSE^E(\eta_r;\ell_r,\gamma), \tag{30}$$

where $AMSE^E(\eta;\ell,\gamma) = \ell + \eta^2 \cdot m \cdot \lambda(\delta\ell;\gamma) - 2\eta \cdot \mu\ell \cdot c^2(\delta\ell;\gamma) \cdot \beta$ *is the AMSE of EBLP* $\hat{S}_i^\eta = \eta \cdot \hat{u}\hat{u}^\top Y_i$ *in a single-spiked model with spike strength* ℓ *under the assumptions of Cor. 2.2. Here* $\lambda(\delta\ell;\gamma)$ *is the limit empirical spike, while* $c^2 = c^2(\delta\ell;\gamma)$ *is the squared cosine, both corresponding to spike strength* $\delta\ell$*, defined in Cor. 2.2. The asymptotically optimal coefficients are*

$$\eta_k^* = \frac{\mu\ell_k c_k^2}{\mu^2\ell_k + m}, \tag{31}$$

where $c_k^2 = c^2(\delta\ell;\gamma)$*.*

The basic discovery is that *the optimal coefficients for empirical PCs are different from those for population PCs.* The coefficients using empirical PCs are reduced by a squared cosine compared to the coefficients using population PCs: $\eta_k^* = c_k^2\tau_k^*$.

Note that we chose the optimal coefficients to minimize the limiting MSE. However, the limiting MSE, and thus the optimal coefficients, depend on the unknown parameters δ, ℓ_k, c_k^2. To make this a practical method, we can estimate the unknown parameters. The missingness parameter δ can be estimated by plug-in. Based on Cor. 2.2, the estimation of ℓ_k and c_k^2 can be done by inverting the spike forward map $\ell_k \to \lambda(\ell_k)$, see e.g., Bai and Ding (2012); Donoho et al. (2013).

It is worth pointing out that the squared error $\|S_i - \hat{S}_i\|^2$ for each individual column S_i of BLP and EBLP does not converge in probability or a.s. In fact, its variability is of unit order, and does not decrease as $n, p \to \infty$. However, as shown in Thm 4.1, its expectation—the MSE—does converge. Furthermore, we will show in Sec. 4.2.2 that the average error of EBLP over the *entire data matrix* converges *almost surely* (to the AMSE for a single column). As we will show, this is because EBLP, when applied to all columns of the data matrix, is a *singular value shrinkage* estimator, defined by modifying the singular values of the data matrix while leaving the empirical singular vectors fixed.

As a consequence, in the case of missing data the optimal AMSE agrees with that achieved by optimal singular value shrinkage described in Nadakuditi (2014) and Gavish and Donoho (2014). We emphasize, however, that Thm. 4.1 applies to *individual* data points, not just to the entire matrix.

4.2.1 Comments on the proof

Part 2 of Thm. 4.1 is a nontrivial result, because the AMSE is determined by stochastically dependent random quantities such as $u_k^\top D_i \hat{u}_k$. These are challenging to study, because D_i and $\hat{u}_k$ are dependent random variables. We will analyze these quantities from first principles.

While we will explain our method in detail later, briefly, we use the *outlier equation* approach (see Lemma 5.14), which reduces studying the inner products $w^\top \hat{u}_k$ to certain inner products $w^\top r$, where $r = r(Y)$ are vectors that depend on the entire dataset but not directly on the singular vector $\hat{u}_k$. As we will see, these inner products are more convenient to study. The basic method was introduced by Benaych-Georges and Nadakuditi (2012), who used it to study the angles between u_k and $\hat{u}_k$. We extend their approach to other angles, which are more challenging to study.

4.2.2 Singular value shrinkage and the almost sure convergence of the error in the reduced-noise model

A well-studied approach to matrix denoising is known as *singular value shrinkage*. Here, the singular values of the data matrix Y are replaced with shrunken versions, analogous to the eigenvalue shrinkage of covariance matrices studied in Sec. 3. When applied to every row of the data matrix Y, EBLP is a singular value shrinkage algorithm: indeed, since $\hat{S}_i^\eta = \sum_{k=1}^r \eta_k \cdot \hat{u}_k \hat{u}_k^\top Y_i$, we can write the entire denoised matrix in the form

$$\hat{S}^\eta = \sum_{k=1}^r \eta_k \cdot \hat{u}_k \hat{u}_k^\top Y = \sum_{k=1}^r \eta_k \sigma_k(Y) \cdot \hat{u}_k \hat{v}_k^\top. \tag{32}$$

In other words, the denoised matrix $\hat{S}^\eta$ has the same singular vectors as the data matrix Y, where the singular values have been moved from $\sigma_k(Y)$ to $\eta_k \sigma_k(Y)$ for $k = 1, \ldots, r$ (and the remaining ones set to 0). It turns out that for the value of η_k^* given by equation (31), $\eta_k \sigma_k(Y)$ is the optimal singular value of the denoised matrix if we seek to minimize the asymptotic Frobenius loss $\|S - \hat{S}^\eta\|_F^2$.

The proof of this fact follows easily from the analysis of the optimal singular value shrinkers given by Gavish and Donoho (2014). This paper derives optimal shrinkers only when there is no missing data, and the task is to recover a low rank matrix from noisy observations of its entries. The k^{th} singular value of the asymptotically optimal shrunken matrix (with respect to Frobenius loss) is $\ell_k^{1/2} c_k \tilde{c}_k$, where c_k is the asymptotic cosine of the angle between the right population singular vector and the right empirical singular vector, and $\tilde{c}_k$ is the asymptotic cosine of the angle between the left population singular vector and the left empirical singular vector. In the models considered in Gavish and Donoho (2014), these values are $c_k = c(\ell_k; \gamma)$, where c is the cosine forward map from (6), and $\tilde{c}_k = \tilde{c}(\ell_k; \gamma) \geq 0$, where $\tilde{c}^2(\ell; \gamma) = (1 - \gamma/\ell^2)/(1 + 1/\ell)$.

As we can see from Cor. 2.2, in our data model the reduced model behaves like the original model, but with spike strengths reduced from ℓ_k to $\delta \ell_k$. In particular, the value of c_k^2 is $c_k^2 = c^2(\delta\ell; \gamma)$. In fact, it is easy to see from the proof of that this is true for the left singular vectors as well; that is, $\tilde{c}^2 = \tilde{c}^2(\delta \ell_k; \gamma)$.

We now quote the result of Gavish and Donoho (2014) that the optimal singular value is equal to $\ell_k^{1/2} c_k \tilde{c}_k$; since formula (32) shows that this is equal to $\eta_k^* \sigma_k(Y)$, the optimal coefficient η_k^* is:

$$\eta_k^* = \frac{\ell_k^{1/2} c(\delta \ell_k; \gamma) \tilde{c}(\delta \ell_k; \gamma)}{\sigma_k(Y)} = \ell_k^{1/2} c^2(\delta \ell_k; \gamma) \frac{\tilde{c}(\delta \ell_k; \gamma)}{\ell_k^{1/2} \sigma_k(Y) c(\delta \ell_k; \gamma)}$$

Substituting the asymptotic value for $\sigma_k(Y)$ given from (7), a straightforward algebraic manipulation shows that

$$\eta_k^* = \ell_k^{1/2} c^2(\delta \ell_k; \gamma) \frac{\ell_k^{1/2}}{\delta \ell_k + 1} = \frac{\ell_k c_k^2}{\delta \ell_k + 1}.$$

This agrees with the value for the optimal coefficient η_k for EBLP derived in Thm. 4.1. In particular, the EBLP estimator we derive for the entire data matrix and the optimal singular value shrinkage estimator are identical. Furthermore, from the analysis of Gavish and Donoho (2014), the error of the *entire* denoised matrix converges *almost surely* to the expression given in equation (30).

We have proved the following theorem:

Theorem 4.2. *If $Y = [Y_1, \ldots, Y_n]^\top$ is the n-by-p matrix of observations in the reduced-noise model, then the asymptotically optimal singular value shrinkage estimator of the n-by-p signal matrix $S = [S_1, \ldots, S_n]^\top$ is equal to the EBLP estimator defined in Thm. 4.1. Furthermore, the squared Frobenius error of this estimator converges a.s. to the formula* (30)*.*

4.2.3 Comparison with matrix completion algorithms

In the case when the reductions matrices are binary, the task of denoising the data matrix Y to approximate S is a *matrix completion* problem – see, for instance, Candès and Recht (2009), Recht (2011), Keshevan et al. (2009) and Keshevan et al. (2010). A typical model for matrix completion is a low-rank matrix with eigenvectors satisfying an incoherence condition, with order $O(n \cdot \mathrm{poly}(\log n))$ entries revealed uniformly at random. In the setting we study in this paper, the number of observed entries of the matrix Y will be $O(n^2)$, almost an order of magnitude more. When the noise level is very small compared to the magnitude of the entries in the clean matrix (or put differently, when $\ell_r \gg 1$), then the smaller number of samples is sufficient to recover the low-rank matrix to high accuracy. The references listed above provide several methods with these guarantees.

We compare the in-sample EBLP denoiser to the OptSpace method for matrix completion found in Keshevan et al. (2009) and Keshevan et al. (2010). Briefly, their method removes rows/columns with too many observations, truncates the singular values of the data matrix Y, and then cleans up the resulting matrix by an iterative algorithm. In Figure 3 we plot the squared Frobenius errors in reconstructing a rank 1 matrix, for different choices of spike size ℓ and missingness parameter δ. In this experiment, both the signal and noise are Gaussian, $\gamma = .8$ and the dimension $p = 400$. Each data point plotted is the average error over 100 Monte Carlo runs of the experiment.

We observe that the in-sample EBLP outperforms OptSpace when ℓ is small relative to the noise level. As the size of ℓ grows, OptSpace's performance improves, and for small δ and large ℓ it outperforms EBLP. This is consistent with the guarantees provided for OptSpace; it does well in the low-noise regime, with a small number of samples.

4.3 Out-of-sample denoising

We now study *out-of-sample* denoising in the reduced-noise model, where we denoise new datapoints from the same distribution using a denoiser constructed on an existing dataset. This is typically faster than recomputing the denoiser on the entire dataset. For the oracle BLP denoiser, which assumes knowledge of u, this is the same as in-sample denoising. For the EBLP, however, it turns out that the optimal shrinkage coefficients in this case are different.

To analyze this case, let $Y_0 = D_0 X_0$ be the new sample from the same distribution. We evaluate the limit of the out-of-sample mean squared prediction error $\mathbb{E}\|S_0 - \hat{S}_0^\eta\|^2$ of the EBLP $\hat{S}_0^\eta = \sum_{k=1}^r \eta_k \hat{u}_k \hat{u}_k^\top Y_0$, where $\hat{u}_k$ were formed based on Y_i, $i \geq 1$.

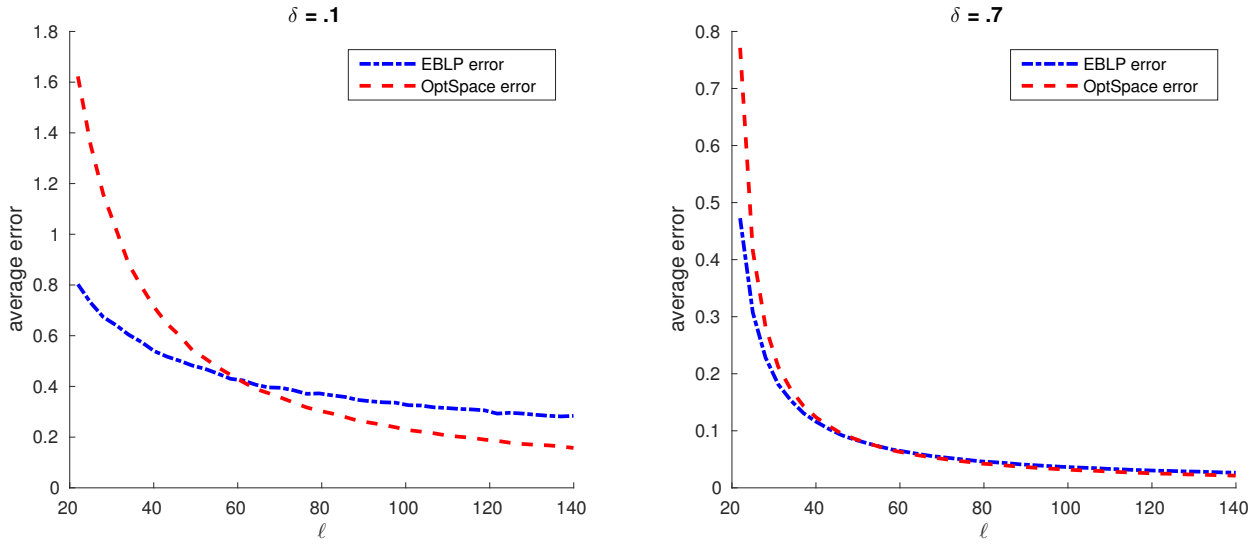

Figure 3: Comparison of EBLP and OptSpace for completing/denoising the matrix, for different signal strengths. Left; $\delta = 1/10$. Right: $\delta = 7/10$. Errors are measured as the squared Frobenius norm between the estimator and the full, clean, rank 1 matrix. Every experiment was averaged over 100 runs, with $\gamma = .8$ and $p = 400$.

Theorem 4.3. *In the setting of Cor. 2.2 consider out-of-sample denoising of a new sample Y_0 using the empirical BLP $\hat{S}_0^{\eta}$ based on the observations Y_i, $i = 1, \ldots, n$. Then the limit of the out-of-sample prediction error is*

$$E^{\eta,o}(\eta_1, \ldots, \eta_r; \ell_1, \ldots, \ell_r, \gamma) = \sum_{k=1}^{r} E^{\eta_k,o}(\eta_r; \ell_r, \gamma),$$

where $E^{\eta,o}(\eta; \ell, \gamma) = \ell + \eta^2 \cdot (m + \mu^2 \ell c^2) - 2\eta \cdot \mu \ell c^2$ is the out-of-sample prediction error of EBLP in a single-spiked model with spike ℓ under the assumptions of Cor. 2.2. The asymptotically optimal shrinkage coefficients are

$$\eta_k^* = \frac{\mu \ell_k c_k^2}{m + \mu^2 \ell_k c_k^2}.$$

See Sec. 5.10 for the proof. The key point is that the optimal shrinkage for out-of-sample prediction is *different* from *both* of the shrinkers from in-sample denoising. We also mention that in the special case when $D_i = I_p$ for all i, the optimal shrinkage formula matches the one obtained by Singer and Wu (2013), under slightly more restrictive assumptions.

It may be counterintuitive that the optimal coefficient changes when a single data point Y_0 is added. However, this can be understood because in high-dimensions, a single data point can drastically change the empirical eigenvectors. For an illustration in a simpler setting, consider the sample covariance matrix based on n samples, $\hat{\Sigma}_{Y,n} = n^{-1} \sum_{k=1}^{n} Y_k Y_k^{\top}$ and the corresponding

sample covariance $\hat{\Sigma}_{Y,n+1}$ based on the $n+1$ samples $Y_0, Y_1, \ldots, Y_n$. Their difference is

$$\hat{\Sigma}_{Y,n} - \hat{\Sigma}_{Y,n+1} = \frac{1}{n} Y_0 Y_0^\top + \frac{1}{n(n+1)} \sum_{k=1}^{n+1} Y_k Y_k^\top = \frac{1}{n} Y_0 Y_0^\top + \frac{1}{n} \hat{\Sigma}_{Y,n+1}.$$

Since the operator norm of $\hat{\Sigma}_{Y,n+1}$ converges a.s. to a finite quantity, the term $\hat{\Sigma}_{Y,n+1}/n$ is asymptotically negligible. However, the operator norm of $Y_0 Y_0^\top/n$ of size $\|Y_0\|^2/n$, which converges a.s. to δ. Since the spectral distributions of $\hat{\Sigma}_{Y,n}$ and $\hat{\Sigma}_{Y,n+1}$ converge to the same value, the fact that $\|\hat{\Sigma}_{Y,n} - \hat{\Sigma}_{Y,n+1}\|_{op}$ is of order 1 is due to the fact that the addition of a single data point Y_0 completely changes the direction of the empirical PCs of the data.

To summarize and better understand our findings, we plot the optimal shrinkage coefficients and MSE for the three scenarios (BLP, in-sample optimal empirical BLP, and out-of-sample optimal empirical BLP) in a single-spiked model with $\gamma = 1/2$ and $D_i = I_p$ for all i, on Fig. 4. The optimal shrinkage for out-of-sample EBLP is intermediate between the stronger in-sample EBLP and the weaker out-of-sample BLP shrinkage coefficients. However, perhaps unexpectedly the MSE for the two EBLP scenarios agrees exactly! This prompts us to state the following result, proved in Sec. 5.11.

Proposition 4.4 (In-sample vs out-of-sample EBLP)**.** *The asymptotically optimal shrinkage coefficient for in-sample EBLP is smaller than the asymptotically optimal shrinkage coefficient for the out-of-sample EBLP. However, the asymptotically optimal MSEs are equal in the two cases.*

This result is interesting, because it shows that same MSE can be achieved out-of-sample as in-sample. Out-of-sample denoising should be harder, because it involves a new datapoint never seen before. The "hardness" of out-of-sample denoising should be observed in the leading order finite sample (n) correction to the asymptotic MSE. However, the above result shows that this correction vanishes as $n, p \to \infty$. Using the right amount of shrinkage, the same MSE can be achieved asymptotically even out of sample.

4.4 Unreduced-noise

We now study the denoising problem under the unreduced-noise model (3), where $Y_i = D_i S_i + \varepsilon_i$ under the assumptions of Cor. 2.2. The analysis is similar to the reduced-noise model (2). The key conclusions are summarized in Table 5.

The BLP of S_i based on Y_i is $\hat{S}_i^{BLP} = \mathrm{Cov}\,[S_i, Y_i]\,\mathrm{Cov}\,[Y_i, Y_i]^{-1}\, Y_i$. Under the conditions of Cor. 2.2, we can show (see Sec. 5.12.1) that this is asymptotically equivalent to $\hat{S}_i^\tau = \sum_{k=1}^r \tau_k u_k u_k^\top Y_i$, where $\tau_k = \mu\ell_k/(\mu^2\ell_k + 1)$.

As before, the AMSE of S_i^τ with arbitrary τ decouples into the AMSEs over the different spikes ℓ_k, and those are equal to the AMSEs for the single-spiked model with spikes equal to ℓ_k. For a single-spiked model with spike ℓ we obtain in Sec. 5.12.1 that

$$E\|S_i - \hat{S}_i^\tau\|^2 \to \ell + \tau^2(\ell\mu^2 + 1) - 2\tau\ell\mu.$$

The optimal coefficient is $\tau^* = \ell\mu/(\ell\mu^2 + 1)$—as it should be, based on the above discussion—and it has an AMSE of $\ell/(\ell\mu^2 + 1)$. The advantage of this calculation is that it provides the MSE for any coefficient τ.

Next, for the EBLP in the multispiked case, we use $\hat{u}_k$ as estimators of u_k. Since the form of the BLP is the same as before, the EBLP scaled by $\eta = (\eta_1, \ldots, \eta_r)$ have the form in (29). To compute

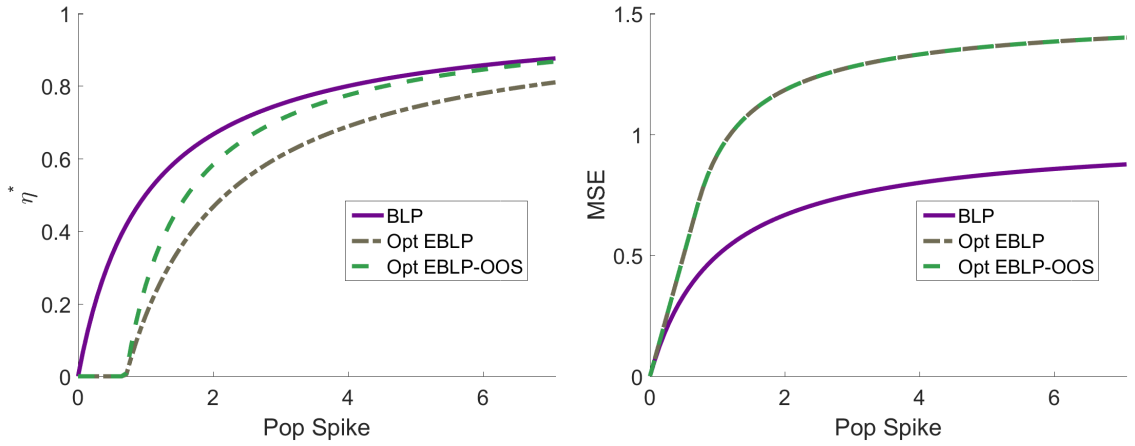

Figure 4: Optimal shrinkage coefficients (left) and MSE (right) for the three scenarios (BLP, in-sample optimal empirical BLP, and out-of-sample optimal empirical BLP) for $\gamma = 1/2$ and $\delta = 1$.

Table 5: Unreduced-noise: Denoising in the single-spiked case. BLP: population Best Linear Predictor using u. EBLP: in-sample empirical Best Linear Predictor using $\hat{u}$. EBLP-OOS: out-of-sample empirical Best Linear Predictor using $\hat{u}$. Here we abbreviate $\lambda = \lambda(\mu^2\ell;\gamma)$, $c^2 = c^2(\mu^2\ell;\gamma)$.

Name	Definition	Asy MSE	Asy Opt η	Asy Opt MSE
BLP	$\eta \cdot uu^\top Y_i$	$\ell + \eta^2 \cdot (\mu^2\ell + 1) - 2\eta \cdot \mu\ell$	$\frac{\mu\ell}{\mu^2\ell+1}$	$\frac{\ell}{\mu^2\ell+1}$
EBLP	$\eta \cdot \hat{u}\hat{u}^\top Y_i$	$\ell + \eta^2 \cdot \lambda - 2\eta \cdot \mu\ell c^2 \cdot [1 + \gamma/(\mu^2\ell)]$	$\frac{\mu\ell c^2}{\mu^2\ell+1}$	$\ell - \lambda\left(\frac{\mu\ell c^2}{\mu^2\ell+1}\right)^2$
EBLP-OOS	$\eta \cdot \hat{u}\hat{u}^\top Y_i$	$\ell + \eta^2 \cdot (\mu^2\ell c^2 + 1) - 2\eta \cdot \mu\ell c^2$	$\frac{\mu\ell c^2}{\mu^2\ell c^2+1}$	$\ell \cdot \frac{\mu^2\ell c^2 s^2+1}{\mu^2\ell c^2+1}$

the AMSE, it is again not hard to see that it decouples into the corresponding single-spiked AMSEs. In the single-spiked case, we find in Sec. 5.12.2 that with $\lambda = \lambda(\mu^2\ell;\gamma)$, $c^2 = c^2(\mu^2\ell;\gamma)$,

$$\mathbb{E}\|S_i - \hat{S}_i^\eta\|^2 \to \ell + \eta^2 \cdot \lambda - 2\eta \cdot \mu\ell c^2 \cdot [1 + \gamma/(\mu^2\ell)].$$

This shows that the optimal coefficient is $\eta^* = \mu\ell c^2/(\mu^2\ell + 1)$.

Finally, for out-of-sample EBLP denoising, it is again not hard to see that the AMSE decouples over the different spikes, and each term equals the AMSE in the single-spiked case. For the AMSE in the single-spiked case, we let (Y_0, D_0) be a new sample, and find (Sec. 5.12.3)

$$\mathbb{E}\|S_0 - \eta\hat{u}\hat{u}^\top Y_0\|^2 \to \ell + \eta^2(1 + \mu^2\ell c^2) - 2\eta\mu\ell c^2.$$

The optimal coefficient is $\eta^* = \mu\ell c^2/(1+\mu^2\ell c^2)$, while the optimal MSE is $\ell(1+\mu^2\ell c^2 s^2)/(1+\mu^2\ell c^2)$. These findings are summarized in Table 5.

4.4.1 Singular value shrinkage and the almost sure convergence of the error in the unreduced-noise model

As we discussed in Sec. 4.2.2 for the reduced-noise model, in-sample EBLP applied to every column of the data matrix $Y = [Y_1, \ldots, Y_n]^\top$ is equal to the asymptotically optimal singular value shrinkage estimator of the clean matrix $S = [S_1, \ldots, S_n]^\top$. The same reasoning applies verbatim to the unreduced-noise model, replacing the almost sure limits of the empirical eigenvalues and angles with their counterparts for unreduced-noise model.

From Cor. 2.2, the value of the asymptotic value of the cosine of the angle between the right empirical singular vector and the right population singular vector is c_k^2 is $c_k^2 = c^2(\mu^2 \ell; \gamma)$; and the same proof of this easily shows that the asymptotic cosine of the angle between the left singular vectors is $\tilde{c}^2 = \tilde{c}^2(\delta \ell_k; \gamma)$.

We now quote the formula from Gavish and Donoho (2014), which says that the optimal singular value is equal to $\ell_k^{1/2} c_k \tilde{c}_k$; since formula (32) shows that this is equal to $\eta_k^* \sigma_k(Y)$, the optimal coefficient η_k^* is:

$$\eta_k^* = \frac{\ell_k^{1/2} c(\delta \ell_k; \gamma) \tilde{c}(\delta \ell_k; \gamma)}{\sigma_k(Y)} = \ell_k^{1/2} c^2(\delta \ell_k; \gamma) \frac{\tilde{c}(\delta \ell_k; \gamma)}{\ell_k^{1/2} \sigma_k(Y) c(\delta \ell_k; \gamma)}$$

Substituting the asymptotic value for $\sigma_k(Y)$ given from Cor. 2.2, it is easy to see that

$$\eta_k^* = \ell_k^{1/2} c^2(\delta \ell_k; \gamma) \frac{\ell_k^{1/2}}{\delta \ell_k + 1} = \frac{\ell_k c_k^2}{\delta \ell_k + 1}.$$

This agrees with the value for the optimal coefficient η_k for EBLP derived in Thm. 4.1. In particular, the EBLP estimator we derive for the entire data matrix and the optimal singular value shrinkage estimator are identical. Furthermore, from the analysis of Gavish and Donoho (2014), the error of the *entire* denoised matrix converges *almost surely* to the expression given in equation (30).

We have proved the following theorem:

Theorem 4.5. *If $Y = [Y_1, \ldots, Y_n]^\top$ is the n-by-p matrix of observations in the unreduced-noise model, then the asymptotically optimal singular value shrinkage estimator of the n-by-p signal matrix $S = [S_1, \ldots, S_n]^\top$ is equal to the EBLP estimator for the unreduced-noise model defined in Sec. 4.4. Furthermore, the squared Frobenius error of this estimator converges a.s. to the AMSE in the unreduced-noise model.*

4.5 Simulations

Next we perform a simulation to check the finite-sample accuracy of our formulas for denoising. In the reduced-noise model with $D_i = I_p$, we consider a single-spiked model with $n = 1000$, $p = 500$, and generate iid Gaussian random variables z_i, ε_i, as well as a standardized iid Gaussian random vector u. We vary the spike strength on a grid, and compare our formulas for in-sample theoretical MSE to those obtained by averaging the denoising error in the first sample $\|\hat{S}_1 - S_1\|_2^2$ over 50 Monte Carlo simulations. The results in Fig. 5 show that the formulas are accurate up to the sampling error. This validates our results from Thm. 4.1.

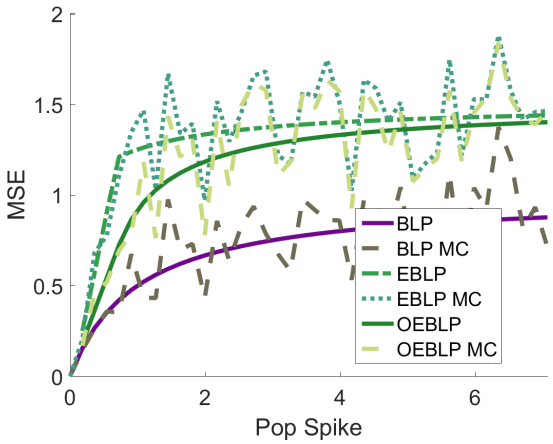

Figure 5: In-sample MSE of denoising schemes: Theoretical results overlaid with Monte Carlo (MC) results. BLP: Best Linear Predictor assuming known population eigenvector u. Emp BLP: Empirical Best Linear Predictor, using the empirical eigenvector $\hat{u}$, with the sub-optimal shrinkage coefficient from the BLP. Opt Emp BLP: Empirical Best Linear Predictor using optimal shrinkage coefficient. No missing data ($\delta = 1$). $\gamma = 1/2$.

Acknowledgements

The authors wish to thank Joakim Anden, Tejal Bhamre, Xiuyuan Cheng, David Donoho, and Iain Johnstone for helpful discussions.

5 Proofs

5.1 Proof of Thm. 2.1

The proof of Thm. 2.1 spans multiple sections, until Sec. 5.2.1. The proof of the claims under observation models (2) and (3) are very similar. Therefore, we present the proof of the result under model (2), and outline the argument for model (3) in Sec. 5.1.5.

Moreover, to illustrate the idea of the proof, we first prove the single-spiked case, i.e., when $r = 1$. The proof of the multispiked extension is provided in Sec. 5.2. The form $D_i = \mu + \Sigma^{1/2} E_i$ of the reduction matrices implies the following decomposition for the observations Y_i:

$$\begin{aligned} Y_i &= (\mu + \Sigma^{1/2} E_i) X_i \\ &= \mu(\ell^{1/2} z_i u + \varepsilon_i) + \Sigma^{1/2} E_i X_i \\ &= \ell^{1/2} z_i \mu u + [\mu \varepsilon_i + \Sigma^{1/2} E_i X_i]. \end{aligned}$$

This suggests a "signal+noise" decomposition for the reduced vectors Y_i. Let us denote by $\nu = \mu u / \xi^{1/2}$ the normalized reduced signal, where $\xi = \|\mu u\|^2 \to \tau$, and by $\varepsilon_i^* = \mu \varepsilon_i + \Sigma^{1/2} E_i X_i$ the noise component. The noise has two parts: $\mu \varepsilon_i$ is due to sampling, while $\Sigma^{1/2} E_i X_i$ is due to projection.

In matrix form, with the $n \times p$ matrix Y having rows $Y_i^\top$:

$$Y = (\xi\ell)^{1/2}\tilde{Z}\nu^\top + \mathcal{E}^*. \tag{33}$$

This suggests that after reduction, the signal strength ℓ changes to $\xi\ell$, while the noise structure changes from ε_i to ε_i^*. This is not obvious, however, because the noise ε_i^* is functionally dependent on the signal X_i. Therefore we cannot rely on existing results. Instead, we will analyze the model from first principles, and show that the dependence is asymptotically negligible. For non-diagonal reduction matrices D_i, the depencence may be asymptotically non-negligible; this explains why we currently need the diagonal assumption.

5.1.1 Proof outline

We will extend the technique of Benaych-Georges and Nadakuditi (2012) to characterize the spiked eigenvalues in the model (33). We denote the normalized vector $Z = n^{-1/2}\tilde{Z}$, the normalized noise $N = n^{-1/2}\mathcal{E}^*$ and the normalized observable matrix $\tilde{Y} = n^{-1/2}Y$. Then, our model is

$$\tilde{Y} = (\xi\ell)^{1/2} \cdot Z\nu^\top + N. \tag{34}$$

We will assume that $n, p \to \infty$ such that $p/n \to \gamma > 0$. For simplicity of notation, we will first assume that $n \le p$, implying that $\gamma \ge 1$. It is easy to see that everything works when $n \ge p$.

By Lemma 4.1 of Benaych-Georges and Nadakuditi (2012), the singular values of $\tilde{Y}$ that are not singular values of N are the positive reals t such that the 2-by-2 matrix

$$M_n(t) = \begin{bmatrix} t \cdot Z^\top(t^2 I_n - NN^\top)^{-1}Z & Z^\top(t^2 I_n - NN^\top)^{-1}N\nu \\ \nu^\top N^\top(t^2 I_n - NN^\top)^{-1}Z & t \cdot \nu^\top(t^2 I_p - N^\top N)^{-1}\nu \end{bmatrix} - \begin{bmatrix} 0 & (\xi\ell)^{-1/2} \\ (\xi\ell)^{-1/2} & 0 \end{bmatrix}$$

is not invertible, i.e., $\det[M_n(t)] = 0$. We will find almost sure limits of the entries of $M_n(t)$, to show that it converges to a deterministic matrix $M(t)$. Solving the equation $\det[M(t)] = 0$ will provide an equation for the almost sure limit of the spiked singular values of $\tilde{Y}$. For this we will prove the following results:

Lemma 5.1 (The noise matrix)**.** *The noise matrix N has the following properties:*

1. *The eigenvalue distribution of $N^\top N$ converges almost surely (a.s.) to the Marchenko-Pastur distribution $F_{\gamma,H}$ with aspect ratio $\gamma \ge 1$.*
2. *The top eigenvalue of $N^\top N$ converges a.s. to the upper edge b_H^2 of the support of $F_{\gamma,H}$.*

This is proved in Sec. 5.1.2. For brevity we write $b = b_H$. Compared to Benaych-Georges and Nadakuditi (2012), the key technical innovation here is to show that the contribution of the reduced signal component $\Sigma^{1/2}E_iS_i$ to $\mathcal{E}^*$ is negligible. This is accomplished by an ad-hoc bound on the operator norm of the contribution.

Since $\tilde{Y}$ is a rank-one perturbation of N, it follows that the eigenvalue distribution of $\tilde{Y}^\top\tilde{Y}$ also converges to the MP law $F_{\gamma,H}$. This proves the first claim of Thm 2.1.

Moreover, since $NN^\top$ has the same n eigenvalues as the nonzero eigenvalues of $N^\top N$, the two facts in Lemma 5.1 imply that when $t > b$, $n^{-1}\operatorname{tr}(t^2 I_n - NN^\top)^{-1} \to \int (t^2 - x)^{-1} d\underline{F}_{\gamma,H}(x) = -\underline{m}(t^2)$. Here $\underline{F}_{\gamma,H}(x) = \gamma F_{\gamma,H}(x) + (1-\gamma)\delta_0$ and $\underline{m} = \underline{m}_{\gamma,H}$ is the Stieltjes transform of $\underline{F}_{\gamma,H}$. Clearly this convergence is uniform in t. As a special note, when t is a singular value of the random matrix N, we formally define $(t^2 I_p - N^\top N)^{-1} = 0$ and $(t^2 I_n - NN^\top)^{-1} = 0$. When $t > b$, the complement of this event happens a.s. In fact, from Lemma 5.1 it follows that $(t^2 I_p - N^\top N)^{-1}$ has a.s. bounded operator norm. Next we control the quadratic forms in the matrix M_n.

Lemma 5.2 (The quadratic forms)**.** *When $t > b$, the quadratic forms in the matrix $M_n(t)$ have the following properties:*

1. $Z^\top (t^2 I_n - NN^\top)^{-1} Z - n^{-1} \operatorname{tr}(t^2 I_n - NN^\top)^{-1} \to 0$ *a.s.*
2. $Z^\top (t^2 I_n - NN^\top)^{-1} N\nu \to 0$ *a.s.*
3. $\nu^\top (t^2 I_p - N^\top N)^{-1} \nu \to -m(t^2)$ *a.s., where $m = m_{\gamma,H}$ is the Stieltjes transform of the Marchenko-Pastur distribution $F_{\gamma,H}$.*

Moreover the convergence of all three terms is uniform in $t > b + c$, for any $c > 0$.

This is proved in Sec. 5.1.3. The key technical innovation is the proof of the third part. Most results for controlling quadratic forms $x^\top A x$ are concentration bounds for random x. Here $x = \nu$ is fixed, and matrix $A = (t^2 I_p - N^\top N)^{-1}$ is random instead. For this reason we adopt the "deterministic equivalents" technique of Bai et al. (2007) for quantities $x^\top (zI_p - N^\top N)^{-1} x$, with the key novelty that we can take the imaginary part of the complex argument to zero. The latter observation is nontrivial, and mirrors similar techniques used recently in universality proofs in random matrix theory (see e.g., the review by Erdős and Yau, 2012).

Lemmas 5.1 and 5.2 will imply that for $t > b$, the limit of $M_n(t)$ is

$$M(t) = \begin{bmatrix} -t \cdot \underline{m}(t^2) & -(\tau\ell)^{-1/2} \\ -(\tau\ell)^{-1/2} & -t \cdot m(t^2) \end{bmatrix}.$$

By the Weyl inequality, $\sigma_2(\tilde{Y}) \le \sigma_2((\xi\ell)^{1/2} \cdot Z\nu^\top) + \sigma_1(N) = \sigma_1(N)$. Since $\sigma_1(N) \to b$ a.s. by Lemma 5.1, we obtain that $\sigma_2(\tilde{Y}) \to b$ a.s. Therefore for any $\varepsilon > 0$, a.s. only $\sigma_1(\tilde{Y})$ can be a singular value of $\tilde{Y}$ in $(b + \varepsilon, \infty)$ that is not a singular value of N.

It is easy to check that $D(x) = x \cdot \underline{m}(x) m(x)$ is strictly decreasing on (b^2, ∞). Hence, denoting $h = \lim_{t \downarrow b} D(t^2)$, for $\tau\ell > h$, the equation $D(t^2) = 1/(\tau\ell)$ has a unique solution $t \in (b, \infty)$. By Lemma A.1 of Benaych-Georges and Nadakuditi (2012), we conclude that for $\tau\ell > h$, $\sigma_1(\tilde{Y}) \to t$ a.s., where t solves the equation $\det[M(t)] = 0$, or equivalently,

$$t^2 \cdot \underline{m}(t^2) m(t^2) = \frac{1}{\tau\ell}.$$

If $\tau\ell \le h$, then we note that $\det[M_n(t)] \to \det[M(t)]$ uniformly on $t > b + \varepsilon$. Therefore, if $\det[M_n(t)]$ had a root $\sigma_1(\tilde{Y})$ in $(b + \varepsilon, \infty)$, $\det[M(t)]$ would also need to have a root there, which is a contradiction. Therefore, we conclude $\sigma_1(\tilde{Y}) \le b + \varepsilon$ a.s., for any $\varepsilon > 0$. Since $\sigma_1(\tilde{Y}) \ge \sigma_2(\tilde{Y}) \to b$, we conclude that $\sigma_1(\tilde{Y}) \to b$ a.s., as desired. This finishes the spike convergence claim in Thm. 2.1.

Next, we turn to proving the convergence of the angles between the population and sample eigenvectors. Let $\hat{Z}$ and $\hat{u}$ be the singular vectors associated with the top singular value $\sigma_1(\tilde{Y})$ of $\tilde{Y}$. Then, by Lemma 5.1 of Benaych-Georges and Nadakuditi (2012), if $\sigma_1(\tilde{Y})$ is not a singular value of X, then the vector $\eta = (\eta_1, \eta_2) = (u^\top \hat{u}, Z^\top \hat{Z})$ belongs to the kernel of the matrix $M_n(\sigma_1(\tilde{Y}))$. By the above discussion, this 2-by-2 matrix is of course singular, so this provides one linear equation for the vector r (with $R = (t^2 I_n - NN^\top)^{-1}$)

$$t\eta_1 \cdot Z^\top R Z + \eta_2 [Z^\top R N \nu - (\xi\ell)^{-1/2}] = 0.$$

By the same lemma cited above, it follows that we have the *norm identity* (with $t = \sigma_1(\tilde{Y})$)

$$t^2 \eta_1^2 \cdot Z^\top R^2 Z + \eta_2^2 \cdot \nu^\top N^\top R^2 N \nu + 2t\eta_1\eta_2 \cdot Z^\top R^2 N \nu = (\xi\ell)^{-1}. \tag{35}$$

This follows from taking the norm of the equation $t\eta_1 \cdot RZ + \eta_2 \cdot RN\nu = (\xi\ell)^{-1/2}\hat{Z}$ (see Lemma 5.1 in Benaych-Georges and Nadakuditi (2012)). We will find the limits of the quadratic forms below.

Lemma 5.3 (More quadratic forms)**.** *The quadratic forms in the norm identity have the following properties:*

1. $Z^\top(t^2 I_n - NN^\top)^{-2}Z - n^{-1}\operatorname{tr}(t^2 I_n - NN^\top)^{-2} \to 0$ *a.s.*
2. $Z^\top(t^2 I_n - NN^\top)^{-2}N\nu \to 0$ *a.s.*
3. $\nu^\top N^\top(t^2 I_n - NN^\top)^{-2}N\nu \to m(t^2) + t^2 m'(t^2)$ *a.s., where* m *is the Stieltjes transform of the Marchenko-Pastur distribution* $F_{\gamma,H}$*.*

The proof is in Sec. 5.1.4. Again, the key novelty is the proof of the third claim. The standard concentration bounds do not apply, because u is non-random. Instead, we use an argument from complex analysis constructing a sequence of functions $f_n(t)$ such that their derivatives are $f_n'(t) = \nu^\top N^\top(t^2 I_n - NN^\top)^{-2}N\nu$, and deducing the convergence of $f_n'(t)$ from that of $f_n(t)$.

Lemma 5.3 implies that $n^{-1}\operatorname{tr}(t^2 I_n - NN^\top)^{-2} \to \int (t^2 - x)^{-2} d\underline{F}_{\gamma,H}(x) = \underline{m}'(t^2)$ for $t > b$. Solving for η_1 in terms of η_2 from the first equation, plugging in to the second, and taking the limit as $n \to \infty$, we obtain that $\eta_2^2 \to c_2$, where

$$c_2\left(\frac{\underline{m}'(t^2)}{\tau\ell\underline{m}(t^2)^2} + m(t^2) + t^2 m'(t^2)\right) = \frac{1}{\tau\ell}.$$

Using $D(x) = x \cdot m(x)\underline{m}(x)$, we find $c_2 = \underline{m}(t^2)/[D'(t^2)\tau\ell]$, where t solves (5). From the first equation, we then obtain $\eta_1^2 \to c_1$, where $c_1 = m(t^2)/[D'(t^2)\tau\ell]$, where t is as above. This finishes the proof of Thm. 2.1 in the single-spiked case. The proof of the multispiked case is a relatively simple extension of the previous argument, so we present it in Sec. 5.2.

5.1.2 Proof of Lemma 5.1

Recall that $N = n^{-1/2}\mathcal{E}^*$, where $\mathcal{E}^*$ has rows $\varepsilon_i^* = \mu\varepsilon_i + \Sigma^{1/2}E_i X_i$. Note

$$\varepsilon_{ij}^* = \mu_j\varepsilon_{ij} + \sigma_j E_{ij}X_{ij} = [\mu_j + \sigma_j E_{ij}]\varepsilon_{ij} + \sigma_j \ell^{1/2}E_{ij}z_i u_j.$$

Since $\varepsilon_i = \Gamma^{1/2}\alpha_i$, the terms $a_{ij} = [\mu_j + \sigma_j E_{ij}]g_j\varepsilon_{ij}$ are independent random variables with variance $g_j^2\mathbb{E}D_{ij}^2 = g_j^2(\mu_j^2 + \sigma_j^2)$. Recall that we assumed that the distribution H_p of $g_j^2\mathbb{E}D_{ij}^2$ converges weakly to the distribution H.

Hence the eigenvalue distribution of the matrix $A^\top A$, where $A = (n^{-1/2}a_{ij})_{ij}$, converges to Marchenko-Pastur distribution $F_{\gamma,H}$ (Bai and Silverstein, 2009, Thm. 4.3). Moreover, since $\mathbb{E}\alpha_{ij}^4 < \infty$ and $\mathbb{E}D_{ij}^4 < \infty$, we have $\mathbb{E}a_{ij}^4 < \infty$. In addition, by assumption $\sup g_j^2\mathbb{E}D_{ij}^2 \to \sup \operatorname{supp}(H)$. Thus the largest eigenvalue of $A^\top A$ converges a.s. to the upper edge b^2 of the support of $F_{\gamma,H}$, see Bai and Silverstein (1998) and (Bai and Silverstein, 2009, Cor. 6.6).

Therefore, since σ_j are bounded, it is enough to show that the operator norm of the error matrix E^{**} with entries $n^{-1/2}E_{ij}z_i u_j$ converges to zero a.s. This will ensure that $N = A + E^{**}$ has the same two properties as A above, namely its ESD and operator norm converge.

Now, denoting by $\odot$ elementwise products

$$\|E^{**}\| = \sup_{\|a\|=\|c\|=1} a^\top E^{**}c = n^{-1/2} \sup_{\|a\|=\|c\|=1} (a \odot z)^\top E(c \odot u).$$

We have $\|a \odot z\| \leq \|a\| \max |z_i| = \max |z_i|$ and $\|c \odot u\| \leq \|c\| \max |u_i| = \max |u_i|$, hence

$$\|E^{**}\| \leq \|E\| \max |z_i| \max |u_i|.$$

Since z has iid standardized entries, and $C := \mathbb{E} z_i^{4+\phi} < \infty$, we can derive that

$$Pr(\max |z_i| \geq a) \leq \mathbb{E} \max |z_i|^{4+\phi} / a^{4+\phi} \leq nC/a^{4+\phi}.$$

Taking $a = n^{1/2-\phi'}$ for ϕ' small enough, we obtain, $\max |z_i| \leq n^{1/2-\phi'}$ a.s.

However, since E_{ij} are iid standardized random variables with bounded 4-th moment, $\|E\| \to 1 + \sqrt{\gamma}$ a.s. (Bai and Silverstein, 2009). Since $\|u\|_\infty \leq C \log(p)^B / p^{1/2}$, we obtain $\|E^{**}\| \leq C(1 + \sqrt{\gamma}) \cdot \log(p)^B n^{1/2-\phi'} p^{-1/2} \to 0$ a.s., as required.

5.1.3 Proof of Lemma 5.2

Since $N = A + E^{**}$, and $\|E^{**}\| \to 0$ a.s., it is enough to show the same concentration statements for A instead of N. Indeed, it is easy to see that the error terms are all negligible.

Part 1: For $Z^\top (t^2 I_n - AA^\top)^{-1} Z$, note that Z has iid entries —with mean 0 and variance $1/n$—that are independent of A. We will use the following result:

Lemma 5.4 (Concentration of quadratic forms, consequence of Lemma B.26 in Bai and Silverstein (2009))**.** *Let $x \in \mathbb{R}^k$ be a random vector with i.i.d. entries and $\mathbb{E}[x] = 0$, for which $\mathbb{E}\left[(\sqrt{k} x_i)^2\right] = 1$ and $\sup_i \mathbb{E}\left[(\sqrt{k} x_i)^{4+\phi}\right] < C$ for some $\phi > 0$ and $C < \infty$. Moreover, let A_k be a sequence of random $k \times k$ symmetric matrices independent of x, with a.s. uniformly bounded eigenvalues. Then the quadratic forms $x^\top A_k x$ concentrate around their means: $x^\top A_k x - k^{-1} \operatorname{tr} A_k \to_{a.s.} 0$.*

We apply this lemma with $x = Z$, $k = p$ and $A_p = (t^2 I_n - AA^\top)^{-1}$. To get almost sure convergence, here it is required that z_i have finite $4 + \phi$-th moment. This shows the concentration of $Z^\top (t^2 I_n - AA^\top)^{-1} Z$.

Part 2: To show $Z^\top (t^2 I_n - AA^\top)^{-1} A\nu$ concentrates around 0, we note that $w = (t^2 I_n - AA^\top)^{-1} A\nu$ is a random vector independent of Z, with a.s. bounded norm. Hence, conditional on w:

$$\begin{aligned} Pr(|Z^\top w| \geq a | w) &\leq a^{-4} \mathbb{E} |Z^\top w|^4 = a^{-4} [\sum_i \mathbb{E} Z_{ni}^4 w_i^4 + \sum_{i \neq j} \mathbb{E} Z_{ni}^2 \mathbb{E} Z_{nj}^2 w_i^2 w_j^2] \\ &\leq a^{-4} \mathbb{E} Z_{n1}^4 (\sum_i w_i^2)^2 = a^{-4} n^{-2} \mathbb{E} Z_1^4 \cdot \|w\|_2^4 \end{aligned}$$

For any C we can write

$$Pr(|Z^\top w| \geq a) \leq Pr(|Z^\top w| \geq a | \|w\| \leq C) + Pr(\|w\| > C).$$

For sufficiently large C, the second term, $Pr(\|w\| > C)$ is summable in n. By the above bound, the first term is summable for any C. Hence, by the Borel-Cantelli lemma, we obtain $|Z^\top w| \to 0$ a.s. This shows the required concentration.

Part 2: Finally we need to show that $\nu^\top (t^2 I_p - A^\top A)^{-1} \nu$ concentrates around a definite value. This is probably the most interesting part, because the vector u is not random. Most results for controlling expressions of the above type are designed for random u; however here the matrix A is random instead. For this reason we will adopt a different approach.

Under our assumption we have $\nu^\top(\Gamma(\Sigma+\mu^2)-zI_p)^{-1}\nu \to m_H(z)$, for $z=t^2+iv$ with $v>0$ fixed. Therefore, Thm 1 of Bai et al. (2007) shows that $\nu^\top(zI_p-A^\top A)^{-1}\nu \to -m(z)$ a.s., where $m(z)$ is the Stieltjes transform of the Marchenko-Pastur distribution $F_{\gamma,H}$.

A close examination of their proofs reveals that their result holds when $v\to 0$ sufficiently slowly, for instance $v=n^{-\alpha}$ for $\alpha=1/10$. The reason is that all bounds in the proof have the rate $N^{-k}v^{-l}$ for some small $k,l>0$, and hence they converge to 0 for v of the above form.

For instance, the very first bounds in the proof of Thm 1 of Bai et al. (2007) are in Eq. (2.2) on page 1543. The first one states a bound of order $O(1/N^r)$. The inequalities leading up to it show that the bound is in fact $O(1/(N^r v^{2r}))$. Similarly, the second inequality, stated with a bound of order $O(1/N^{r/2})$ is in fact $O(1/(N^{r/2}v^r))$. These bounds go to zero when $v=n^{-\alpha}$ with small $\alpha>0$. In a similar way, the remaining bounds in the theorem have the same property.

To get the convergence for real t^2 from the convergence for complex $z=t^2+iv$, we note that

$$|\nu^\top(zI_p-A^\top A)^{-1}\nu-\nu^\top(t^2I_p-A^\top A)^{-1}\nu| = v|\nu^\top(zI_p-A^\top A)^{-1}(t^2I_p-A^\top A)^{-1}\nu| \leq$$
$$\leq v\|(t^2I_p-A^\top A)^{-1}\|^2\cdot u^\top u.$$

As discussed above, when $t>b$, the matrices $(t^2I_p-A^\top A)^{-1}$ have a.s. bounded operator norm. Hence, we conclude that if $v\to 0$, then a.s.

$$\nu^\top(zI_p-A^\top A)^{-1}\nu-\nu^\top(t^2I_p-A^\top A)^{-1}\nu\to 0.$$

Finally, $m(z)\to m(t^2)$ by the continuity of the Stieltjes transform for all $t^2>0$ (Bai and Silverstein, 2009). We conclude that $\nu^\top(t^2I_p-A^\top A)^{-1}\nu\to -m(t^2)$ a.s. This finishes the analysis of the last quadratic form.

5.1.4 Proof of Lemma 5.3

As in Lemma 5.2, it is enough to show the same concentration statements for A instead of N.

Parts 1 and 2: The proof of Part 1 and 2 are exactly analogous to those in Lemma 5.2. Indeed, the same arguments work despite the change from $(t^2I_p-N^\top N)^{-1}$ to $(t^2I_p-N^\top N)^{-2}$, because the only properties we used are its independence from Z, and its a.s. bounded operator norm. These also hold for $(t^2I_p-N^\top N)^{-2}$, so the same proof works.

Part 3: We start with the identity $\nu^\top N^\top(t^2I_n-NN^\top)^{-2}N\nu = -\nu^\top(t^2I_p-N^\top N)^{-1}\nu + t^2\nu^\top(t^2I_p-N^\top N)^{-2}u$. Since in Lemma 5.2 we have already established $\nu^\top(t^2I_p-N^\top N)^{-1}\nu\to -m(t^2)$, we only need to show the convergence of $\nu^\top(t^2I_p-N^\top N)^{-2}u$.

For this we will employ the following *derivative trick* (see e.g., Dobriban and Wager, 2015). We will construct a function with two properties: (1) its derivative is the quantity $\nu^\top(t^2I_p-N^\top N)^{-2}u$ that we want, and (2) its limit is convenient to obtain. The following lemma will allow us to get our answer by interchanging the order of limits:

Lemma 5.5 (see Lemma 2.14 in Bai and Silverstein (2009)). *Let $f_1,f_2,\ldots$ be analytic on a domain D in the complex plane, satisfying $|f_n(z)|\leq M$ for every n and z in D. Suppose that there is an analytic function f on D such that $f_n(z)\to f(z)$ for all $z\in D$. Then it also holds that $f_n'(z)\to f'(z)$ for all $z\in D$.*

Accordingly, consider the function $f_p(r)=-\nu^\top(rI_p-N^\top N)^{-1}\nu$. Its derivative is $f_p'(r)=\nu^\top(rI_p-N^\top N)^{-2}u$. Let $\mathcal{S}:=\{x+iv: x>b+\varepsilon\}$ for a sufficiently small $\varepsilon>0$, and let us work on

the set of full measure where $\|N^\top N\| < b+\varepsilon/2$ eventually, and where $f_p(r) \to m(r)$. By inspection, f_p are analytic functions on $\mathcal{S}$ bounded as $|f_p| \le 2/\varepsilon$. Hence, by Lemma 5.5, $f_p'(r) \to m'(r)$.

In conclusion, $\nu^\top N^\top (t^2 I_p - N^\top N)^{-2} N\nu \to m(t^2) + t^2 m'(t^2)$, finishing the proof.

5.1.5 Proof of Thm. 2.1: Model (3)

The proof for model (3) is very similar to that for model (2). Therefore, we only present the outline. Working again in the single-spiked case for simplicity, we have the following decomposition for the observations Y_i:

$$\begin{aligned} Y_i &= (\mu + \Sigma^{1/2} E_i) S_i + \varepsilon_i \\ &= \ell^{1/2} z_i \mu u + [\mu E_i S_i + \varepsilon_i]. \end{aligned}$$

Denoting $\nu = \mu u/\xi^{1/2}$, where $\xi = \|\mu u\|^2 \to \tau$, and $\varepsilon_i^* = \mu E_i S_i + \varepsilon_i$, we have in matrix form

$$Y = (\xi\ell)^{1/2} \tilde{Z} \nu^\top + \mathcal{E}^*.$$

As in the proof of Lemma 5.1, it is not hard to see that the operator norm $\|\mathcal{E}^* - \mathcal{E}\|_{\text{op}} \to 0$. Therefore, the spectral properties of Y are equivalent to those of $\tilde{Y} = (\xi\ell)^{1/2} \tilde{Z} \nu^\top + \mathcal{E}$. However, this is now a spiked model where the signal component is independent of the noise component.

It follows immediately from Thm 4.3 of Bai and Silverstein (2009) that the singular value distribution of $\mathcal{E} = [\alpha_1, \dots, \alpha_n]^\top \Gamma^{1/2}$ converges to the general Marchenko-Pastur distribution $F_{\gamma,G}$, where G is the limit of the distributions G_p of g_j^2, $j = 1, \dots, p$. Similarly the top singular value of $\mathcal{E}$ converges to the upper edge b_G of $F_{\gamma,G}$.

Moreover, it also follows that the analogues of Lemmas 5.2 and 5.3 hold in our case. Indeed, the same arguments carry through, because the same assumptions hold. This allows the entire argument from Sec. 5.1.1 to carry through, finishing the single-spiked case of Thm. 2.1. The extension to the multispiked case is analogous to that in model (2).

5.1.6 Numerical computation of the quantities from Thm. 2.1

SPECTRODE computes the Marchenko-Pastur forward map: given an input limit population spectrum H and an aspect ratio γ, it outputs an accurate numerical approximation to the limit empirical spectral distribution (ESD) $F_{\gamma,H}$. Dobriban (2015) established the numerical convergence of the method, and showed in experiments that it is much faster than previous proposals. The method is publicly available at `http://github.com/dobriban/eigenedge`.

The output of SPECTRODE includes a numerical approximation $\hat{m}$ to the Stieltjes transform m of the limit ESD, computed over a dense grid x_i on the real line. It also includes an approximation $\hat{b}^2$ of the upper edge b^2 of the ESD. From this, we compute an approximation of the D-transform as $\hat{D}(x_i) = x_i \cdot \hat{m}(x_i) \cdot \hat{\underline{m}}(x_i)$, where $\hat{\underline{m}}(x_i) = \gamma \cdot \hat{m}(x_i) + (\gamma - 1)/x_i$. Since D is monotone decreasing on (b^2, ∞), we find the smallest grid point x_i such that $\hat{D}(x_i) \le 1/\ell$ to approximately compute $D^{-1}(1/\ell)$.

Finally, the derivative $\underline{m}'(x)$ can be expressed as a function $\underline{m}'(x) = \mathcal{F}(\underline{m}(x))$ by differentiating the Marchenko-Pastur fixed-point equation (see e.g., Dobriban, 2015). Therefore, we compute a numerical approximation to $D'(x) = m(x) \cdot \underline{m}(x) + x[m(x) \cdot \underline{m}'(x) + m'(x) \cdot \underline{m}(x)]$ by approximating $\hat{\underline{m}}'(x)$ via the same function $\hat{\underline{m}}'(x) = \mathcal{F}(\hat{\underline{m}}(x))$. Similarly we approximate $\hat{m}'$. With these steps, we obtain a full numerical implementation of Thm. 2.1.

5.2 Proof of Thm. 2.1 - Multispiked extension

Let us denote by $u_i = \mu u_i/\xi_i^{1/2}$ the normalized reduced signals, where $\xi_i = \|\mu u_i\|^2 \to \tau_i$, and by $\varepsilon_i^* = \mu\varepsilon_i + \Sigma^{1/2}E_iX_i$. For the proof we start as in Sec. 5.1.1, obtaining $Y_i = \sum_{k=1}^r (\xi_k\ell_k)^{1/2} z_{ik}\nu_k + \varepsilon_i^*$. Defining the $r \times r$ diagonal matrices L, Δ with diagonal entries ℓ_k, ξ_k (respectively), and the $n \times r$, $p \times r$ matrices $Z, \mathcal{V}$, with columns $Z_k = n^{-1/2}(z_{1k}, \ldots, z_{nk})^\top$ and u_k respectively, we have

$$\tilde{Y} = Z(\Delta L)^{1/2}U^\top + \mathcal{E}^*.$$

The matrix $M_n(t)$ is now $2r \times 2r$, and has the form

$$M_n(t) = \begin{bmatrix} t \cdot Z^\top(t^2 I_n - NN^\top)^{-1}Z & Z^\top(t^2 I_n - NN^\top)^{-1}N\mathcal{V} \\ \mathcal{V}^\top N^\top(t^2 I_n - NN^\top)^{-1}Z & t \cdot \mathcal{V}^\top(t^2 I_p - N^\top N)^{-1}\mathcal{V} \end{bmatrix} - \begin{bmatrix} 0_r & (\Delta L)^{-1/2} \\ (\Delta L)^{-1/2} & 0_r \end{bmatrix}.$$

It is easy to see that Lemma 5.1 still holds in this case. To find the limits of the entries of M_n, we need the following additional statement.

Lemma 5.6 (Multispiked quadratic forms)**.** *The quadratic forms in the multispiked case have the following properties for $t > b$:*

1. $Z_k^\top R^\alpha Z_j \to 0$ *a.s. for* $\alpha = 1, 2$, *if* $k \neq j$.
2. $\nu_k^\top (t^2 I_p - N^\top N)^{-\alpha}\nu_j \to 0$ *a.s. for* $\alpha = 1, 2$, *if* $k \neq j$.

This lemma is proved in Sec. 5.2.1, using similar techniques as those in Lemma 5.1. Defining the $r \times r$ diagonal matrices T with diagonal entries τ_k, we conclude that for $t > b$, $M_n(t) \to M(t)$ a.s., where now

$$M(t) = \begin{bmatrix} -t \cdot \underline{m}(t^2) I_r & -(TL)^{-1/2} \\ -(TL)^{-1/2} & -t \cdot m(t^2) I_r \end{bmatrix}.$$

As before, by Lemma A.1 of Benaych-Georges and Nadakuditi (2012), we get that for $\tau_k\ell_k > 1/D(b^2)$, $\sigma_k(\tilde{Y}) \to t_k$ a.s., where $t_k^2 \cdot \underline{m}(t_k^2)m(t_k^2) = 1/(\tau_k\ell_k)$. This finishes the spike convergence proof.

To obtain the limit of the angles for $\hat{u}_k$ for a k such that $\ell_k > \tau_k D(b^2)$, consider the left singular vectors $\hat{Z}_k$ associated to $\sigma_k(\tilde{Y})$. Define the $2r$-vector

$$\alpha = \begin{bmatrix} \beta_1 \\ \beta_2 \end{bmatrix} = \begin{bmatrix} (\Delta L)^{1/2}\mathcal{V}^\top \hat{u}_k \\ (\Delta L)^{1/2} Z^\top \hat{Z}_k \end{bmatrix}.$$

The vector α belongs to the kernel of $M_n(\sigma_k(\tilde{Y}))$. As argued by Benaych-Georges and Nadakuditi (2012), the fact that the projection of α into the orthogonal complement of $M(t_k)$ tends to zero, implies that $\alpha_j \to 0$ for all $j \notin \{k, k+r\}$. This proves that $\nu_j^\top \hat{u}_k \to 0$ for $j \neq k$, and the analogous claim for the left singular vectors.

The linear equation $M_n(\sigma_k(\tilde{Y}))\alpha = 0$ in the k-th coordinate, where $k \leq r$, reads (with $t = \sigma_k(\tilde{Y})$):

$$t\alpha_k Z_k^\top R Z_k - \alpha_{r+k}(\xi_k\ell_k)^{-1/2} + \sum_{i \neq k} M_n(\sigma_k(\tilde{Y}))_{ik}\alpha_k = 0.$$

Only the first two terms are non-negligible due to the behavior of M_n, so we obtain $t\alpha_k Z_k^\top R Z_k = \alpha_{r+k}(\xi_k \ell_k)^{-1/2} + o_p(1)$. Moreover taking the norm of the equation $\hat{Z}_k = R(tZ\beta_1 + N\mathcal{V}\beta_2)$ (see Lemma 5.1 in Benaych-Georges and Nadakuditi (2012)), we get

$$t^2 \sum_{i,j \leq r} \alpha_i \alpha_j Z_i^\top R^2 Z_j + \sum_{i,j \leq r} \alpha_{k+i}\alpha_{k+j} \nu_i^\top N^\top R^2 N \nu_j + \sum_{i,j \leq r} \alpha_i \alpha_{k+j} Z_i R^2 N \nu_j = 1.$$

From Lemma 5.6 and the discussion above, only the terms $\alpha_k^2 Z_k^\top R^2 Z_k$ and $\alpha_{r+k}^2 \nu_k^\top N^\top R^2 N \nu_k$ are non-negligible, so we obtain

$$t^2 \alpha_k^2 Z_k^\top R^2 Z_k + \alpha_{r+k}^2 \nu_k^\top N^\top R^2 N \nu_k = 1 + o_p(1).$$

Combining the two equations above,

$$\alpha_{r+k}^2 \left[\frac{Z_k^\top R^2 Z_k}{\xi_k \ell_k (Z_k^\top R Z_k)^2} + \nu_k^\top N^\top R^2 N \nu_k \right] = 1 + o_p(1).$$

Since this is the same equation as in the single-spiked case, we can take the limit in a completely analogous way. This finishes the proof.

5.2.1 Proof of Lemma 5.6

As in Lemma 5.2, it is enough to show the same concentration statements for A instead of N.

Part 1: The convergence $Z_k^\top R^\alpha Z_j \to 0$ a.s. for $\alpha = 1, 2$, if $k \neq j$, follows directly from the following well-known lemma, cited from Couillet and Debbah (2011):

Lemma 5.7 (Proposition 4.1 in Couillet and Debbah (2011))**.** *Let $x_n \in \mathbb{R}^n$ and $y_n \in \mathbb{R}^n$ be independent sequences of random vectors, such that for each n the coordinates of x_n and y_n are independent random variables. Moreover, suppose that the coordinates of x_n are identically distributed with mean 0, variance C/n for some $C > 0$ and fourth moment of order $1/n^2$. Suppose the same conditions hold for y_n, where the distribution of the coordinates of y_n can be different from those of x_n. Let A_n be a sequence of $n \times n$ random matrices such that $\|A_n\|$ is uniformly bounded. Then $x_n^\top A_n y_n \to_{a.s.} 0$.*

Part 2: To show $\nu_k^\top (t^2 I_p - N^\top N)^{-\alpha} \nu_j \to 0$ a.s. for $\alpha = 1, 2$, if $k \neq j$, the same technique cannot be used, because the vectors u_k are deterministic. However, it is straightforward to check that the method of Bai et al. (2007) that we adapted in proving Part 3 of Lemma 5.2 extends to proving $\nu_k^\top (t^2 I_p - N^\top N)^{-1} \nu_j \to 0$. Indeed, it is easy to see that all their bounds hold unchanged. In the final step, as a deterministic equivalent for $\nu_k^\top (t^2 I_p - N^\top N)^{-1} \nu_j \to 0$, one obtains $\nu_k^\top (t^2 I_p - t^2 m(t^2)\Sigma)^{-1} \nu_j$, which tends to 0 by our assumption, showing $\nu_k^\top (t^2 I_p - N^\top N)^{-1} \nu_j \to 0$. Then $\nu_k^\top (t^2 I_p - N^\top N)^{-2} \nu_j \to 0$ follows from the derivative trick employed in Part 3 of Lemma 5.3. This finishes the proof.

5.3 Proof of Corollary 2.2

This corollary is a special case of our previous results, so the convergence results hold in this case. We only need to check that the limits are given by the formulas provided. Since very similar analysis has been performed by Benaych-Georges and Nadakuditi (2012) and Nadakuditi (2014), we only give part of the proof.

Under model (2), since we consider the singular values of the normalized matrix $m^{-1/2}\tilde{Y}$ instead of $\tilde{Y}$ as in Thm 2.1, it is easy to see that the relevant equation for the limiting singular values in this case is $t^2 \cdot m_0(t^2)\underline{m}_0(t^2) = 1/\delta\ell$ instead of $t^2 \cdot m(t^2)\underline{m}(t^2) = 1/\tau\ell$.

Note that $\underline{m}_0(t^2) = \int (x-t^2)^{-1} d\underline{F}_{\gamma,H}(x) = \gamma \cdot m_0(t^2) + (\gamma-1)t^{-2}$. Thus the equation for t^2 reads $(\gamma t^2 m_0(t^2)+\gamma-1)m_0(t^2) = (\delta\ell)^{-1}$. However, it is well-known that $m_0(t^2)$ obeys the equation $\gamma t^2 m_0^2 + (t^2+\gamma-1)m_0 + 1 = 0$. From these two relations we obtain $t^2 \cdot m_0(t^2) = -1 - 1/(\delta\ell)$. Plugging this back into the second equation and simplifying, we obtain $t^2 = (\delta\ell+1)(1+\gamma/(\delta\ell))$, as required. Finally, under model (3), the proof is analogous.

5.4 Proofs for covariance estimation

5.4.1 Proof of Cor. 3.2

From the spectral analysis of Y contained in Corollary 2.2 from Section 2 and the formula (10), the proof of this corollary is immediate from the first part of the following lemma, proved below:

Lemma 5.8. *We have the following limits (in operator norm) of the diagonals:*

1. $\lim_{n\to\infty} \|\operatorname{diag}(\hat{\Sigma}_Y) - m \cdot I_p\| = 0$ *a.s.*
2. $\lim_{n\to\infty} \|\operatorname{diag}(\hat{\Sigma}_S) - I_p\| = 0$ *a.s.*

First, note that the second statement is an immediate corollary of the first. In the proof of the first statement, for notational convenience only, we will assume that $r = 1$; however, an identical proof goes through for any fixed r. We will therefore drop the subscript k for the proof.

Decompose X_i into signal and noise, writing $X_i = S_i + \varepsilon_i$, where $S_i = \ell^{1/2} z_i u$ and $\operatorname{Cov}(\varepsilon) = I_p$, and S_i and ε_i are independent. So if $S_i = (s_{i1},\dots,s_{ip})^\top$, $X_i = (X_{i1},\dots,X_{ip})^\top$, $Y_i = (Y_{i1},\dots,Y_{ip})^\top$, and $\varepsilon_i = (\varepsilon_{i1},\dots,\varepsilon_{ip})^\top$, we can write $Y_{ij} = D_{ij}X_{ij} = D_{ij}s_{ij} + D_{ij}\varepsilon_{ij}$ and consequently $Y_{ij}^2 = D_{ij}^2 s_{ij}^2 + 2D_{ij}^2 s_{ij}\varepsilon_{ij} + D_{ij}^2\varepsilon_{ij}^2$.

The $(i,i)^{th}$ element of $\operatorname{diag}(\hat{\Sigma}_Y)$ is then

$$\frac{1}{n}\sum_{j=1}^n Y_{ij}^2 = \frac{1}{n}\sum_{j=1}^n D_{ij}^2 s_{ij}^2 + \frac{1}{n}\sum_{j=1}^n 2D_{ij}^2 s_{ij}\varepsilon_{ij} + \frac{1}{n}\sum_{j=1}^n D_{ij}^2\varepsilon_{ij}^2$$

and we will control each of the three sums on the right side separately.

Observe that $\mathbb{E}\left[\frac{1}{n}\sum_{j=1}^n D_{ij}^2 s_{ij}^2\right] = m\ell u_i^2$ and so

$$\operatorname{Var}\left(\frac{1}{n}\sum_{j=1}^n D_{ij}^2 s_{ij}^2\right) = \frac{1}{n}(\mathbb{E}[d^4]\mathbb{E}[z^4]\ell^2 u_i^4 - m^2\ell u_i^4) \le c\frac{\log^{4B}(n)}{n^3}$$

where $c > 0$ is a constant. Chebyshev's inequality then gives

$$\mathbb{P}\left\{\left|\frac{1}{n}\sum_{j=1}^n D_{ij}^2 s_{ij}^2 - m\cdot\ell^2\cdot u_i^2\right| \ge \varepsilon \text{ for some } i\right\} \le c\frac{p\log^{4B}(n)}{\varepsilon^2 n^3}.$$

Since $u_i^2 \le C\log^B(p)/p$, this shows that $n^{-1}\sum_{j=1}^n D_{ij}^2 s_{ij}^2$ converges a.s. to 0 as $p,n\to\infty$ and $p/n\to\gamma$.

For the sum of terms $D_{ij}^2\varepsilon_{ij}^2$, observe that $\mathbb{E}n^{-1}\sum_{j=1}^n 2D_{ij}^2 s_{ij}\varepsilon_{ij} = 0$, and

$$\mathbb{E}|2D_{ij}^2 s_{ij}\varepsilon_{ij}|^4 = 16\mathbb{E}(d^8)\ell^2 u_i^4 \leq c\frac{\log^{4B}(n)}{n^2}$$

where $c > 0$ is a constant and we have used the estimate $|u_i| \leq C\log^B(p)/\sqrt{p}$. By using Markov's inequality for the fourth moment (see Petrov (2012)), we then obtain:

$$\mathbb{P}\left\{\left|\frac{1}{n}\sum_{j=1}^n 2D_{ij}^2 s_{ij}\varepsilon_{ij}\right| \geq \varepsilon \text{ for some } i\right\} \leq c\frac{\log^{4B}(n)}{\varepsilon^4 n^2}.$$

This proves that $n^{-1}\sum_{j=1}^n 2D_{ij}^2 s_{ij}\varepsilon_{ij}$ converges a.s. to 0 as $p, n \to \infty$ and $p/n \to \gamma$.

Finally, to deal with $n^{-1}\sum_{j=1}^n D_{ij}^2\varepsilon_{ij}^2$, observe that $\mathbb{E}D_{ij}^2\varepsilon_{ij}^2 = m$, and that $\mathbb{E}|D_{ij}^2\varepsilon_{ij}^2|^4 = \mathbb{E}d^8\mathbb{E}\varepsilon^8$ which is assumed finite. Therefore using again Markov's inequality for the fourth moment,

$$\mathbb{P}\left\{\left|\frac{1}{n}\sum_{j=1}^n D_{ij}^2\varepsilon_{ij}^2 - m\right| \geq \varepsilon \text{ for some } i\right\} \leq cp\frac{\mathbb{E}d^8\mathbb{E}\varepsilon^8}{\varepsilon^4 n^4}.$$

This proves that $n^{-1}\sum_{j=1}^n D_{ij}^2\varepsilon_{ij}^2$ converges a.s. to 0 as $p, n \to \infty$ and $p/n \to \gamma$. This finishes the proof.

5.4.2 Proof of (13)

The following analysis is adapted from that in Donoho et al. (2013). We start with a fact from linear algebra.

Lemma 5.9. *Suppose A and B are two p-by-p symmetric matrices, with eigenvectors $u_1, \ldots, u_p$ and $v_1, \ldots, v_p$. Suppose that each one has eigenvalue 0 with multiplicity $p-r$ for some $1 \leq r < p/2$, with corresponding eigenvectors $u_{r+1}, \ldots, u_p$ and $v_{r+1}, \ldots, v_p$. Suppose that no eigenvector $v_1, \ldots, v_r$ lies in the span of $u_1, \ldots, u_r$. Then there is an orthogonal matrix W such that*

$$W^\top AW = \mathrm{diag}(\ell_1, \ldots, \ell_r) \oplus \mathbf{0}_{(p-r)\times(p-r)} \tag{36}$$

and

$$W^\top BW = \tilde{B} \oplus \mathbf{0}_{(p-2r)\times(p-2r)} \tag{37}$$

where $\tilde{B}$ is a $2r$-by-$2r$ matrix whose entries are continuous functions of the inner products $u_i^\top v_j$, $1 \leq i, j \leq r$, on $\mathbb{R}^{2r} \setminus \{\pm 1\}^r$.

Furthermore, if we assume in addition that $u_i^\top v_j = c_i\delta_{i=j}$ for some numbers $c_i \in (-1, 1)$, and we let $s_i = \sqrt{1-c_i^2}$, then the matrix $\tilde{B}$ is of the form:

$$\begin{bmatrix} \mathrm{diag}(\eta_1 c_1^2, \ldots, \eta_r c_r^2) & \mathrm{diag}(\eta_1 c_1 s_1, \ldots, \eta_r c_r s_r) \\ \mathrm{diag}(\eta_1 c_1 s_1, \ldots, \eta_r c_r s_r) & \mathrm{diag}(\eta_1 s_1^2, \ldots, \eta_r s_r^2) \end{bmatrix}$$

Proof. Form vectors $\tilde{v}_1, \ldots, \tilde{v}_{p-2r}$, by performing Gram-Schmidt orthogonalization of $v_1, \ldots, \tilde{v}_{p-2r}$ against the vectors $u_1, \ldots, u_r$. If the columns of W are the vectors in this basis, then clearly (36) holds, since $Au_k = \ell_k u_k$ for $k = 1, \ldots, r$, and $A\tilde{v}_k = 0$ for $k = 1, \ldots, p-2r$. Furthermore, since

$B\tilde{v}_k = 0$ for $k = r+1, \ldots, p-2r$, (37) holds for some $2r$-by-$2r$ block $\tilde{B}$. We need only check that the entries of $\tilde{B}$ are continuous functions of $u_i^\top v_j$, $1 \le i, j \le r$, on $\mathbb{R}^{2r} \setminus \{\pm 1\}^r$.

Let W_{2r} denote the p-by-$2r$ matrix containing the first $2r$ columns of W, U_r denote the matrix with columns $u_1, \ldots, u_r$, V_r denote the matrix with columns $v_1, \ldots, v_r$, and $\tilde{V}_r$ the matrix with columns $\tilde{v}_1, \ldots, \tilde{v}_r$. We can therefore write $W_{2r} = [U_r \;\; \tilde{V}_r]$;

By the Gram-Schmidt construction, every vector v_k, $k = 1, \ldots, r$, can be written as a linear combination of the vectors $u_1, \ldots, u_r, \tilde{v}_1, \ldots, \tilde{v}_k$; in matrix form, these linear combinations can be expressed as:

$$V_r = U_r(U_r^\top V_r) + \tilde{V}_r(\tilde{V}^\top V_r)$$

Now, we observe that the $2r$-by-$2r$ block $\tilde{B}$ is

$$W_{2r}^\top B W_{2r} = \begin{bmatrix} U_r^\top V \\ \tilde{V}_r^\top V \end{bmatrix} \operatorname{diag}(\eta_1, \ldots, \eta_r)[U_r^\top V \;\; \tilde{V}_r^\top V]$$

We need only show that the entries of $\tilde{V}_r^\top V_r$ are continuous functions of the inner products $u_i^\top v_j$, $1 \le i, j \le r$. Let $R = \tilde{V}_r^\top V$. Then $R_{k,l} = 0$ whenever $k > l$. When $k = l$, we have:

$$R_{k,k} = \tilde{v}_k^\top v_k = \sqrt{1 - \sum_{i=1}^{r} (u_i^\top v_k)^2 - \sum_{i=1}^{k-1} (\tilde{v}_i^\top v_k)^2}$$

and when $k < l$ we have

$$R_{k,l} = \frac{-\sum_{i=1}^{r} (u_i^\top v_l)(u_i^\top v_k) + \sum_{i=1}^{k-1} (\tilde{v}_i^\top v_l)(\tilde{v}_i^\top v_k)}{\|v_k - \sum_{i=1}^{r} (u_i^\top v_k) u_k\|}.$$

Since for any $l = 2, \ldots, r$, we have

$$R_{1,l} = \tilde{v}_1^\top v_l = \frac{-\sum_{i=1}^{r} (u_i^\top v_1)(u_i^\top v_l)}{\|v_1 - \sum_{i=1}^{r} (u_i^\top v_1) u_i\|}$$

which is obviously a continuous function of the inner products $v_1^\top u_i$ on $\mathbb{R}^r \setminus \{\pm 1\}^r$, an induction argument easily shows all entries of R are continuous in the inner products $u_k^\top v_l$.

Finally, if $u_i^\top v_j = c_i \delta_{i=j}$, the final assertion follows immediately, finishing the proof. □

If we apply the permutation $\pi_{2r} = (1, r+1, 2, r+2, \ldots, r, 2r)$ to the rows and columns of the block matrix $\tilde{B}$ from Lemma 5.9, the corresponding matrix becomes $\bigoplus_{i=1}^{r} B_2(\eta_i, c_i, s_i)$ where

$$B_2(\eta, c, s) = \begin{pmatrix} \eta c^2 & \eta cs \\ \eta cs & \eta s^2 \end{pmatrix}. \tag{38}$$

Applying the same permutation to the top $2r$-by-$2r$ block of the matrix A from Lemma 5.9 turns this block into: $\bigoplus_{i=1}^{r} A_2(\ell_i)$ where $A_2(\ell) = \begin{pmatrix} \ell & 0 \\ 0 & 0 \end{pmatrix}$.

Now, let $\tilde{B}$ be given by the formula (38), but where $c = c_i = c(\delta \ell_i)$ is the asymptotic inner product between the population eigenvector u_i and the empirical eigenvector $\hat{u}_i$, and $\eta = \eta_i = \eta(\lambda_i)$ is the almost sure limit of the i^{th} eigenvalue $\eta(\hat{\lambda}_i)$ of $\hat{\Sigma}_S^\eta$. Since the shrinkers η collapse the vicinity

of the bulk to 0, it follows from Cor. 3.1 that the shrunken estimator has rank at most r a.s. Hence, by Lemma 5.9 and Cor. 3.1, if $L_p(A, B)$ is any orthogonally-invariant loss function that decomposes over blocks, we have the almost sure convergence of $L_p(\Sigma_S, \hat{\Sigma}_S^\eta)$ to the quantity

$$L_\infty(\Sigma_S, \hat{\Sigma}_S^\eta) = L_{2r}\Big(\bigoplus_{k=1}^{r} A_2(\ell_k), \bigoplus_{k=1}^{r} B_2(\eta(\lambda_k), c_k, s_k)\Big).$$

This is the desired result.

5.4.3 Proof of Props. 3.3

We will only provide a proof for the reduced-noise model; the proof for the unreduced-model is even simpler. Define the operator $\mathcal{L}_p$ acting on p-by-p matrices by $\mathcal{L}_p(\Sigma) = \mu^2\Sigma + \sigma^2 \operatorname{diag}(\Sigma)$ and define the operator L_p by $L_p(\Sigma) = \frac{1}{n}\sum_{k=1}^n D_k \Sigma D_k$. Also, define define the p-by-p matrices $B_p = \hat{\Sigma}_Y - \frac{1}{n}\sum_{k=1}^n D_k^2$ and $\mathcal{B}_p = \hat{\Sigma}_Y - mI_p$.

Let $\Delta B_p = \mathcal{B}_p - B_p$, and $\Delta L_p = \mathcal{L}_p - L_p$. The operator ΔL_p taking p-by-p matrices to p-by-p matrices is diagonal. In this proof, when we refer to the $(i,j)^{th}$ diagonal entry of a diagonal operator on $\mathbb{R}^{p\times p}$, we mean the entry that multiplies the $(i,j)^{th}$ coordinate of a matrix. When $i \neq j$, the $(i,j)^{th}$ diagonal entry of $\Delta L_p(\Sigma)$ is $\mu^2 - n^{-1}\sum_{k=1}^n D_{i,k}D_{j,k}\Sigma_{ij}$; and when $i = j$, the $(i,i)^{th}$ diagonal entry of ΔL_p is $m - n^{-1}\sum_{k=1}^n D_{i,k}^2$.

A proof nearly identical to the proof of Cor. 3.1 shows that the maximum of all the p^2 diagonal entries of ΔL_p converges almost surely to 0; in other words, if $\mathbb{R}^{p\times p}$ is equipped with Frobenius norm, then the operator norm of $\Delta L_p : \mathbb{R}^{p\times p} \to \mathbb{R}^{p\times p}$ converges to 0 almost surely.

The p-by-p matrix ΔB_p is diagonal, with i^{th} diagonal entry equal to $m - n^{-1}\sum_{k=1}^n D_{k,i}$. Again, a proof like the proof of Cor. 3.1 shows that the supremum of these elements converges to zero almost surely as $n, p \to \infty$.

Also, note that the Frobenius norm of the matrix $\hat{\Sigma}_S$, which is the sum of squares of its eigenvalues, is of size $\approx \sqrt{n}$. To see this, observe that

$$\frac{1}{n}\|\hat{\Sigma}_S\|_F^2 = \frac{1}{n}\sum_{k=1}^{p} \sigma_k(\hat{\Sigma}_S)^2 = \frac{1}{n}\sum_{k=1}^{p-r} \sigma_k(\hat{\Sigma}_S)^2 + \frac{1}{n}\sum_{k=p-r+1}^{p} \sigma_k(\hat{\Sigma}_S)^2.$$

The first term $\frac{1}{n}\sum_{k=1}^{p-r}\sigma_k(\hat{\Sigma}_S)^2$ converges almost surely to the second moment of the Marchenko-Pastur law, which is finite since the distribution has finite support. The second term $\frac{1}{n}\sum_{k=p-r+1}^{p}\sigma_k(\hat{\Sigma}_S)^2$ converges to 0.

Observe that

$$L_p(\hat{\Sigma}_S) - \Delta L_p(\hat{\Sigma}_S) = \mathcal{L}_p(\hat{\Sigma}_S) = \mathcal{B}_p = B_p + \Delta B_p = L_p(\hat{\Sigma}'_S) + \Delta B_p$$

or in other words,

$$\frac{\hat{\Sigma}_S - \hat{\Sigma}'_S}{\|\hat{\Sigma}_S\|} = \frac{L_p^{-1}(\Delta L_p \hat{\Sigma}_S)}{\|\hat{\Sigma}_S\|} - \frac{L_p^{-1}(\Delta B_p)}{\|\hat{\Sigma}_S\|}.$$

The result follows immediately.

5.5 Proof that BLP is asymptotically diagonalized in PC basis (Sec. 4.1)

First, it is easy to check that $\mathrm{Cov}\,[S_i, Y_i] = \mu \sum_{k=1}^r \ell_k u_k u_k^\top$, and $M = \mathrm{Cov}\,[Y_i, Y_i]^{-1} = [\mu^2 \sum_{j=1}^r \ell_j u_j u_j^\top + mI_p + E]^{-1}$, where $E = \sigma^2 \sum_{j=1}^r \ell_j \,\mathrm{diag}(u_j \odot u_j)$. Therefore the BLP equals $\hat{S}_i^{BLP} = \mu \sum_{k=1}^r \ell_k u_k u_k^\top M Y_i$. Moreover, since u_j are delocalized, the operator norm $\|E\| \to 0$, so that it is easy to check that $\|\hat{S}_i^{BLP} - \hat{S}_i^0\|^2 \to 0$, where

$$\hat{S}_i^0 = \mu \sum_{k=1}^r \ell_k u_k u_k^\top M_0 Y_i,$$

and $M_0 = [\mu^2 \sum_{j=1}^r \ell_j u_j u_j^\top + mI_p]^{-1}$. Therefore, it is enough to show that $\|\hat{S}_i^0 - \hat{S}_i\|^2 \to 0$, where recall that $\hat{S}_i = \sum_{k=1}^r \mu\ell_k/(\mu^2\ell_k + m) u_k u_k^\top Y_i$. We can write, with $m_k = u_k^\top (M_0 - I_p/(\mu^2\ell_k + m)) Y_i$.

$$\begin{aligned}\|\hat{S}_i^0 - \hat{S}_i\|^2 &= \|\mu \sum_{k=1}^r \ell_k u_k u_k^\top M_0 Y_i - \sum_{k=1}^r \mu\ell_k/(\mu^2\ell_k + m) u_k u_k^\top Y_i\|^2 \\ &= \|\mu \sum_{k=1}^r \ell_k u_k m_k\|^2 = \mu^2 \sum_{k,j=1}^r \ell_k \ell_j m_k m_j u_k^\top u_j.\end{aligned}$$

Therefore, to show $\|\hat{S}_i^0 - \hat{S}_i\|^2 \to 0$, it is enough to show $m_k \to 0$. For this, using the formula $u^\top (uu^\top + T)^{-1} = u^\top T^{-1}/(1 + u^\top T^{-1} u)$, we can write

$$u_k^\top \left[\mu^2 \sum_{j=1}^r \ell_j u_j u_j^\top + mI_p\right]^{-1} = u_k^\top \left[\mu^2 \sum_{j\neq k}^r \ell_j u_j u_j^\top + mI_p\right]^{-1} \Big/ \left(1 + \mu^2 u_k^\top \left[\mu^2 \sum_{j\neq k}^r \ell_j u_j u_j^\top + mI_p\right]^{-1} u_k\right).$$

Under the assumptions of Cor. 2.2, we have $u_k^\top u_j \to \delta_{kj}$, hence it is easy to see that the denominator converges a.s. to $1 + \mu^2/m$. Indeed by using the formula $(VV^\top + mI)^{-1} = [I - V(V^\top V + mI)^{-1} V^\top]/m$, for $V = \mu[\ell_1^{1/2} u_1, \ldots, \ell_r^{1/2} u_r]$ (excluding u_k)

$$u_k^\top \left[\mu^2 \sum_{j\neq k}^r \ell_j u_j u_j^\top + mI\right]^{-1} u_k = u_k^\top u_k/m - u_k^\top V(V^\top V + mI)^{-1} V^\top u_k/m.$$

Now the entries of the $r-1$-dimensional vector $v_k = V^\top u_k$ are $\mu \ell^{1/2} u_j^\top u_k$ for $j \neq k$. Therefore, they converge to zero a.s. Since the operator norm of $(V^\top V + mI)^{-1}$ is bounded above by $1/m$, this shows that the second term converges to zero a.s. This shows that the denominator converges to $1 + \mu^2/m$ a.s. Finally, by using the formula $(VV^\top + mI)^{-1} = [I - V(V^\top V + mI)^{-1} V^\top]/m$ once again, we conclude similarly that

$$u_k^\top \left[\mu^2 \sum_{j\neq k}^r \ell_j u_j u_j^\top + mI\right]^{-1} Y_i = u_k^\top Y_i/m - u_k^\top V(V^\top V + mI)^{-1} V^\top Y_i/m.$$

Denoting $M_k = \left[\mu^2 \sum_{j\neq k}^r \ell_j u_j u_j^\top + I\right]^{-1}$, we have

$$m_k = \frac{u_k^\top M_k Y_i}{1 + u_k^\top M_k u_k} - \frac{u_k^\top Y_i/m}{1 + \mu^2/m \cdot \ell_k} = \frac{u_k^\top Y_i}{m}\left(\frac{1}{1 + u_k^\top M_k u_k} - \frac{1}{1 + \mu^2/m \cdot \ell_k}\right) - \frac{v_k^\top (V^\top V + mI)^{-1} v_k/m}{1 + u_k^\top M_k u_k}.$$

Based on our previous calculations, both terms tend to 0, implying $m_k \to 0$. Therefore, we have $\|\hat{S}_i^{BLP} - \hat{S}_i\|^2 \to 0$. By the triangle inequality, this immediately implies that all MSE properties for $\hat{S}_i^{BLP}$ also hold for $\hat{S}_i$. This proves the desired claim.

5.6 Proof of Thm. 4.1

5.6.1 BLP

For the BLP, consider first the single-spiked case where the squared prediction error of $\hat{S}_i^\tau = \hat{S}_i^{\tau,B} = \tau u u^\top Y_i$, with $Y_i = D_i(S_i + \varepsilon_i) = D_i(\ell^{1/2} z_i u + \varepsilon_i)$ is

$$\begin{aligned}\|S_i - \hat{S}_i^\tau\|^2 &= \|\ell^{1/2} z_i u - \tau u^\top Y_i u\|^2 = (\ell^{1/2} z_i - \tau u^\top Y_i)^2 = [\ell^{1/2} z_i - \tau(\ell^{1/2} z_i u^\top D_i u + u^\top D_i \varepsilon_i)]^2 \\ &= (1 - \tau u^\top D_i u)^2 \ell z_i^2 + \tau^2 (u^\top D_i \varepsilon_i)^2 - 2\tau(1 - \tau u^\top D_i \varepsilon_i)\ell^{1/2} z_i u^\top D_i \varepsilon_i.\end{aligned}$$

Now, z_i and ε_i are independent of D_i, hence, conditional on D_i we have

$$\mathbb{E}[\|S_i - \hat{S}_i^\tau\|^2 | D_i] = (1 - \tau u^\top D_i u)^2 \ell + \tau^2 u^\top D_i^2 u.$$

Moreover, $\mathbb{E} u^\top D_i u = \mu$, $\mathbb{E}(u^\top D_i u)^2 = \mu^2 + \sigma^2 \|u\|_4^4 \to \mu^2$—since u is delocalized—and $\mathbb{E} u^\top D_i^2 u = m$, so the overall expectation converges to $\mathbb{E}\|S_i - \hat{S}_i^\tau\|^2 \to (1 - 2\tau\mu + \tau^2\mu^2)\ell + \tau^2 m$.

This proves that the asymptotic MSE of BLP is $AMSE^B(\tau; \ell, \gamma) = (1 - \tau\mu)^2 \ell + \tau^2 m$. The minimum is achieved for $\tau^* = \mu\ell/(\mu^2\ell + m)$, and equals $m\ell/(\mu^2\ell + m)$.

In the multispiked case, the prediction error is

$$\begin{aligned}\|S_i - \hat{S}_i^\tau\|^2 &= \|\sum_{k=1}^r \ell_k^{1/2} z_{ik} u_k - \tau_k u_k^\top Y_i u_k\|^2 \\ &= \sum_{k=1}^r (\ell_k^{1/2} z_{ik} - \tau_k u_k^\top Y_i)^2 + \sum_{k \neq j} (\ell_k^{1/2} z_{ik} - \tau_k u_k^\top Y_i)(\ell_j^{1/2} z_{ij} - \tau_j u_j^\top Y_i) \cdot u_k^\top u_j\end{aligned}$$

Now, z_{ik}, z_{ij} are independent if $k \neq j$, hence we have if $k \neq j$

$$\mathbb{E}(\ell_k^{1/2} z_{ik} - \tau_k u_k^\top Y_i)(\ell_j^{1/2} z_{ij} - \tau_j u_j^\top Y_i) = \tau_k \tau_j \mathbb{E} u_k^\top Y_i u_j^\top Y_i.$$

Using $Y_i = D_i X_i$, $X_i = \sum_{k=1}^r \ell_k^{1/2} z_{ik} u_k + \varepsilon_i$, and the independence properties described above, we have

$$\begin{aligned}&\mathbb{E} u_k^\top Y_i u_j^\top Y_i = \mathbb{E} u_k^\top D_i X_i X_i^\top D_i u_j = \mathbb{E} u_k^\top D_i (\sum_{m=1}^r \ell_m u_m u_m^\top + I_p) D_i u_j \\ &= \sum_{m=1}^r \ell_m \cdot \mathbb{E}(u_k^\top D_i u_m \cdot u_j^\top D_i u_m) + \mathbb{E} u_k^\top D_i^2 u_j \\ &= \sum_{m=1}^r \ell_m \cdot [\mu^2 (u_k^\top u_m) \cdot (u_j^\top u_m) + \sigma^2 (u_k \odot u_m)^\top \cdot (u_j \odot u_m)] + m u_k^\top u_j.\end{aligned}$$

Above we denoted by $c = a \odot b$ the vector with entries $c_j = a_j b_j$. Since the vectors u_k are delocalized and we assumed $u_k^\top u_l \to 0$ for $k \neq l$, it is easy to see that the entire expression converges to zero.

This shows that

$$\mathbb{E}\|S_i - \hat{S}_i^\tau\|^2 - \mathbb{E}\sum_{k=1}^{r}(\ell_k^{1/2} z_{ik} - \tau_k u_k^\top Y_i)^2 \to 0.$$

Therefore, the asymptotic MSE decouples over the different PCs. Therefore, we can use our previous results about the asymptotic MSE in the single-spiked model, as soon as we can show that the same formulas for the MSE of specific coordinates given in Sec. 5.6.1 hold in this setting. However, this is easy to see using a calculation similar to the one given above. Indeed, we have

$$\begin{aligned}\mathbb{E}(\ell_k^{1/2} z_{ik} - \tau_k u_k^\top Y_i)^2 &= \mathbb{E}(\ell_k^{1/2}(1 - \tau_k u_k^\top D_i u_k) z_{ik} - \sum_{m\neq k} \tau_k \ell_m^{1/2} z_{im} u_k^\top D_i u_m - \tau_k u_k^\top D_i \varepsilon_i)^2 \\ &= \ell_k \mathbb{E}(1 - \tau_k u_k^\top D_i u_k)^2 + \tau_k^2 \sum_{m\neq k} \ell_m \mathbb{E}(u_k^\top D_i u_m)^2 + \tau_k^2 \mathbb{E} u_k^\top D_i^2 u_k.\end{aligned}$$

For $k \neq m$, the term $\mathbb{E}(u_k^\top D_i u_m)^2 = \mu^2 (u_k^\top u_m)^2 + \sigma^2 \|u_k \odot u_m\|^2 \to 0$, while the other terms can be evaluated as before, leading to the same formulas. Therefore, the asymptotic limit of the MSE is the same as that in the single-spiked case. This finishes the claims about BLP.

5.6.2 Empirical BLP

The squared error of a general EBLP denoiser is

$$\begin{aligned}\|S_i - \hat{S}_i^\eta\|^2 &= \|\sum_{k=1}^{r} \ell_k^{1/2} z_{ik} u_k - \sum_{k=1}^{r} \eta_k \hat{u}_k \hat{u}_k^\top Y_i\|^2 \\ &= \sum_{k=1}^{r} \|\ell_k^{1/2} z_{ik} u_k - \eta_k \hat{u}_k \hat{u}_k^\top Y_i\|^2 + \sum_{k\neq j} (\ell_k^{1/2} z_{ik} u_k - \eta_k \hat{u}_k \hat{u}_k^\top Y_i)^\top (\ell_j^{1/2} z_{ij} u_j - \eta_j \hat{u}_j \hat{u}_j^\top Y_i)\end{aligned}$$

The following lemma, proved in Sec. 5.6.3, shows that the cross terms vanish:

Lemma 5.10 (Vanishing cross term in MSE)**.** *The cross terms in the MSE of EBLP denoisers vanish, i.e., for all $k \neq j$, $\mathbb{E}(\ell_k^{1/2} z_{ik} u_k - \eta_k \hat{u}_k \hat{u}_k^\top Y_i)^\top (\ell_j^{1/2} z_{ij} u_j - \eta_j \hat{u}_j \hat{u}_j^\top Y_i) \to 0$*

Therefore, the limit of the MSE in the multispiked case decouples into the MSEs for the individual spikes:

$$\mathbb{E}\|S_i - \hat{S}_i^\eta\|^2 - \sum_{k=1}^{r} \mathbb{E}\|\ell_k^{1/2} z_{ik} u_k - \eta_k \hat{u}_k \hat{u}_k^\top Y_i\|^2 \to 0.$$

To evaluate the limiting MSE, we rely on the following lemma (proved in Sec. 5.7), which finds limiting expectations of inner products of the empirical singular vector $\hat{u}_B$ with the samples Y_i and the population singular vector u.

Lemma 5.11 (Denoising risk limit)**.** *We have the following convergences:*

1. $\mathbb{E}(\hat{u}_k^\top Y_i)^2 \to m\lambda_k^2$, *where* $\lambda_k^2 = t^2(\delta\ell_k; \gamma)$ *is defined in Eq.* (7)*.*
2. $\mathbb{E} z_i (\hat{u}_k^\top Y_i)(u_k^\top \hat{u}_k) \to \mu \ell_k^{1/2} c_k^2 \cdot \beta_k$, *where* $c_k^2 = c^2(\delta\ell_k; \gamma)$ *is defined in Eq.* (8)*, and*

$$\beta = 1 + \frac{\gamma}{\delta \ell_k}. \tag{39}$$

The key technical innovation is the proof of the second part. The main argument is an extension of the technique introduced by Benaych-Georges and Nadakuditi (2012) for characterizing the limits of the inner products $u_j^\top \hat{u}_k$ between population and sample eigenvectors. For our proof, we need to extend this technique to characterize limits $w^\top \hat{u}_k$ for *arbitrary random vectors* w. Since this technique relies on an equation for the "outliers" among the eigenvalues of a finite-rank perturbation of a matrix in the terminology commonly used in random matrix theory (see e.g., Tao, 2013), we call this the *outlier equation* method.

Applying the outlier equation method is nontrivial in our case, because the random vectors w to which we need to apply it—e.g., X_i—are dependent with $\hat{u}_k$. For this reason, we need to use rank-one perturbation formulas once again to show that the dependence is negligible. We envision that the proof method could have several other applications.

Going back to our main argument, based on Lemma 5.11, the limit of the MSE is the following deterministic quantity: $AMSE = \sum_{k=1}^{r} \ell_k + \eta_k^2 \cdot m t_k^2 - 2\eta_k \cdot \mu \ell_k c_k^2 \cdot \beta_k$. The optimal η minimizing the AMSE has $\eta_k^* = \mu \ell_k c_k^2 \cdot \beta_k / [m t_k^2] = \mu \ell_k c_k^2 / [m(1 + \delta \ell_k)]$. This finishes the EBLP analysis.

5.6.3 Proof of Lemma 5.10

We need to show that $\mathbb{E}(\ell_k^{1/2} z_{ik} u_k - \eta_k \hat{u}_k \hat{u}_k^\top Y_i)^\top (\ell_j^{1/2} z_{ij} u_j - \eta_j \hat{u}_j \hat{u}_j^\top Y_i) \to 0$. We expand the parantheses and note that the first term is $\mathbb{E}(\ell_k^{1/2} z_{ik} u_k)^\top \ell_j^{1/2} z_{ij} u_j = 0$, while the last term is a multiple of $\mathbb{E}\hat{u}_k^\top \hat{u}_j \cdot \hat{u}_k^\top Y_i \cdot \hat{u}_j^\top Y_i = 0$ (because $\hat{u}_k^\top \hat{u}_j = 0$ for $k \neq j$). Thus it is enough to show the following claim for $k \neq j$:

$$\mathbb{E} z_{ij} \hat{u}_k^\top u_j \cdot \hat{u}_k^\top Y_i \to 0$$

For this, we have $u_k^\top \hat{u}_j \to 0$ a.s., by Cor. 2.2, thus also $z_{ij} u_k^\top \hat{u}_j \to 0$, so by convergence reduction (Lemma 5.12) it is enough to show that $\mathbb{E}(\hat{u}_k^\top Y_i)^2$ is uniformly bounded. However, as in the proof of Part 1 of Lemma 5.11 in Sec. 5.7.1, we obtain that $\mathbb{E}(\hat{u}_k^\top Y_i)^2$ is uniformly bounded. This finishes the proof of Lemma 5.10.

5.7 Proof of Lemma 5.11

For simplicity of exposition, we first prove the single-spiked case, when $r = 1$. The extension to the multispiked case is presented in Sec. 5.9.

5.7.1 Part 1

Similar to Lee et al. (2010) in the unreduced case, we use the following exchangeability argument:

$$\mathbb{E}(\hat{u}^\top Y_i)^2 = \mathbb{E}\hat{u}^\top Y_i Y_i^\top \hat{u} = n^{-1} \sum_{i=1}^{n} \mathbb{E}\hat{u}^\top Y_i Y_i^\top \hat{u} = n^{-1} \mathbb{E}\hat{u}^\top Y^\top Y \hat{u} = \mathbb{E}\sigma_1(\tilde{Y})^2.$$

Now $\sigma_1(\tilde{Y})^2 \to m \cdot \lambda(\delta\ell; \gamma)$ a.s. by Cor. 2.2. Moreover, by a bound on the expectation of top eigenvalues of sample covariance matrices such as that in Srivastava and Vershynin (2013), it is easy to see that $\mathbb{E}\sigma_1(\tilde{Y})^2$ is uniformly bounded. Hence, $\mathbb{E}\sigma_1(\tilde{Y})^2 \to m \cdot \lambda(\delta\ell; \gamma)$, finishing the proof.

5.7.2 Part 2

As a preliminary remark for Part 2, denoting $t^2 = \lambda(\delta\ell;\gamma)$ we note that $\beta = 1+\gamma/(\delta\ell) = 1 - m\cdot\gamma m(t^2)/[1+m\cdot\gamma m(t^2)]$. By a simple calculation, this is equivalent to $m\cdot m(t^2) = -1/(\gamma+\delta\ell)$. This can be checked as in the proof of Cor. 2.2 in Sec. 5.3.

Therefore it is enough to show that

$$\mathbb{E}z_i(\hat{u}^\top Y_i)(u^\top\hat{u}) \to \ell^{1/2}c^2\left[\mu + \frac{-m\cdot\gamma m(t^2)}{1+m\cdot\gamma m(t^2)}\right] = \mu\ell^{1/2}c^2 + \ell^{1/2}c^2\cdot\frac{-m\cdot\gamma m(t^2)}{1+m\cdot\gamma m(t^2)}.$$

Expanding by using $Y_i = D_iX_i$, $X_i = \ell^{1/2}z_iu+\varepsilon_i$, we have

$$\mathbb{E}z_i(\hat{u}^\top Y_i)(u^\top\hat{u}) = \ell^{1/2}\mathbb{E}z_i^2(\hat{u}^\top D_iu)(u^\top\hat{u}) + \mathbb{E}z_i(\hat{u}^\top D_i\varepsilon_i)(u^\top\hat{u}). \tag{40}$$

Therefore, it is enough to show that $\mathbb{E}z_i^2(\hat{u}^\top D_iu)(u^\top\hat{u}) \to \mu c^2$ and $\mathbb{E}z_i(\hat{u}^\top D_i\varepsilon_i)(u^\top\hat{u}) \to \ell^{1/2}\cdot c^2\cdot[-m\cdot\gamma m(t^2)]/[1+m\cdot\gamma m(t^2)]$. Since we already know that $(u^\top\hat{u})^2 \to c^2$ a.s., our first step is to reduce these two claims by "getting rid" of the $u^\top\hat{u}$ term. For this and similar arguments, we will rely repeatedly on the following simple lemma:

Lemma 5.12 (Convergence Reduction)**.** *Suppose W_n,Y_n are random variables such that $\mathbb{E}W_n \to D$ for a constant D, $\mathbb{E}W_n^2$ is uniformly bounded, $Y_n \to C$ a.s. for some constant C, and Y_n is a.s. uniformly bounded. Then $\mathbb{E}W_nY_n \to DC$.*

In the special case when $C = 0$, the statement $\mathbb{E}W_n \to D$ is not needed.

Proof. We have $\mathbb{E}W_nY_n = \mathbb{E}W_n(Y_n - C) + C\mathbb{E}W_n$, so it is enough to show $\mathbb{E}W_n(Y_n - C) \to 0$. Therefore we can assume without loss of generality that $C = 0$. In that case,

$$|\mathbb{E}W_nY_n| \le |\mathbb{E}W_n^2|^{1/2}|\mathbb{E}Y_n^2|^{1/2} \le M_W|\mathbb{E}Y_n^2|^{1/2} \to 0,$$

because $\mathbb{E}W_n^2$ are uniformly bounded; and because $Y_n \to 0$ a.s. and Y_n are bounded, so $\mathbb{E}(Y_n)^2 \to 0$ by the dominated convergence theorem. Clearly, once we assumed $C = 0$ we did not use $\mathbb{E}W_n \to D$. This finishes the proof. □

To use this lemma, let us choose the orientation of $\hat{u}$ such that $u^\top\hat{u} \ge 0$. Therefore, we know from our main results (e.g., Thm. 2.1), that $u^\top\hat{u} \to c$ a.s. Moreover, since $\|\hat{u}\| = 1$, it is easy to see that we can apply Lemma 5.12 to both terms in (40). Specifically, for the first term, we apply it with $W_n = z_i^2(\hat{u}^\top D_iu)$ and $Y_n = u^\top\hat{u}$; while for the second term we apply it with $W_n = z_i(\hat{u}^\top D_i\varepsilon_i)$ and $Y_n = u^\top\hat{u}$. As mentioned above, this effectively removes the $u^\top\hat{u}$ terms, and thus simplifies the analysis considerably. Indeed, by Lemma 5.12 we conclude that Part 2 follows from the following auxiliary convergence results, proved in Sec. 5.8:

Lemma 5.13 (Denoising risk auxiliary limits: Part 2)**.** *We have the following convergence results:*

1. $\mathbb{E}z_i^2(\hat{u}^\top D_iu) \to \mu c$, *where c^2 is defined in Eq.* (8)*.*
2. $\mathbb{E}z_i(\hat{u}^\top D_i\varepsilon_i) \to \ell^{1/2}\cdot c\cdot\frac{-m\cdot\gamma m(t^2)}{1+m\cdot\gamma m(t^2)}$, *where t^2 is defined in Eq.* (7)*, and m is the Stieltjes transform of the standard Marchenko-Pastur law.*

Based on the above lemma, we completed the proof of Part 2 of Lemma 5.11. We will prove Part 1 of Lemma 5.11 later in Sec. 5.7.1, because it re-uses many of the techniques and results established in the proof of Part 1. This finishes the proof of Lemma 5.11.

5.8 Proof of Lemma 5.13

In this lemma, we need to understand the asymptotics of inner products such as $u^\top D_i \hat{u}$. Previous results by Benaych-Georges and Nadakuditi (2012) characterized the asymptotics of the cosines $u^\top \hat{u}$. However, these results do not allow us to understand the asymptotics of the inner products $u^\top D_i \hat{u}$, due to the dependence of $\hat{u}$ and D_i. Instead, we must go back to first principles, and extend the *outlier equation* technique introduced by Benaych-Georges and Nadakuditi (2012) to our setting. We will see that the current case is more challenging than the one handled in Benaych-Georges and Nadakuditi (2012).

Let us recall the notation from Sec. 5.1.1. According to (34), the normalized data matrix $m^{-1/2}\tilde{Y} = (nm)^{-1/2}Y$ can be written as $m^{-1/2}\tilde{Y} = \ell^{1/2} \cdot Z u^\top + N$. Note here that $u = \nu$. Let here $t = t_p$ be the singular value of $m^{-1/2}\tilde{Y}$ with left and right singular vectors $\hat{Z}$, $\hat{u}$, and suppose that t is not a singular value of N. As in Lemmas 4.1 and 5.1 of Benaych-Georges and Nadakuditi (2012), we then have the following *outlier equation* that provides an equation for the singular vectors of the perturbed matrix $m^{-1/2}\tilde{Y}$:

$$\begin{bmatrix} t \cdot (t^2 I_n - NN^\top)^{-1} & (t^2 I_n - NN^\top)^{-1} N \\ N^\top (t^2 I_n - NN^\top)^{-1} & t \cdot (t^2 I_p - N^\top N)^{-1} \end{bmatrix} \begin{bmatrix} (u^\top \hat{u}) \cdot Z \\ (Z^\top \hat{Z}) \cdot u \end{bmatrix} = \frac{1}{\ell^{1/2}} \begin{bmatrix} \hat{Z} \\ \hat{u} \end{bmatrix} \tag{41}$$

The idea is to *take inner products* of the right hand side with the quantity we want to characterize (e.g., inner product with $D_i u$ to understand $u^\top D_i \hat{u}$), and then evaluate the limit of the quantities on the left hand side. This extends the technique of Benaych-Georges and Nadakuditi (2012), who took inner products only with the population and sample singular vectors (u and $\hat{u}$). In contrast, by using other vectors, such as $D_i u$ and $D_i \varepsilon_i$, we are able to extend vastly the reach of their technique.

To formalize this, let w be a p-dimensional vector. By taking the inner product of the outlier equation's last p coordinates with w, we obtain the following scalar equation:

$$(u^\top \hat{u}) \cdot w^\top N^\top (t^2 I_n - NN^\top)^{-1} Z + t(Z^\top \hat{Z}) \cdot w^\top (t^2 I_p - N^\top N)^{-1} u = \frac{1}{\ell^{1/2}} w^\top \hat{u}. \tag{42}$$

Therefore, we can formalize the outlier equation technique as follows:

Lemma 5.14 (Outlier equation method)**.** *Under the assumptions of Cor. 2.2, suppose that $\hat{u}$, $\hat{Z}$ are chosen so that $\hat{u}^\top u \geq 0$, $\hat{Z}^\top Z \geq 0$. Suppose moreover that $w = w_p$ is a sequence of random vectors such that the assumptions of Lemma 5.12 for W_n apply to $W_{n1} = w^\top N^\top (t^2 I_n - NN^\top)^{-1} Z$ and $W_{n2} = w^\top (t^2 I_p - N^\top N)^{-1} u$. Specifically, suppose that $\mathbb{E}W_{n1} \to w_1^*$ and $\mathbb{E}W_{n2} \to w_2^*$. Then the random variables $w^\top \hat{u}$ also converge in expectation, namely*

$$\mathbb{E} w^\top \hat{u} \to \ell^{1/2} (c(\delta\ell; \gamma) w_1^* + t \cdot \tilde{c}(\delta\ell; \gamma) w_2^*). \tag{43}$$

Here $c \geq 0$ where c^2 is defined in Eq. (8), while $\tilde{c}(\delta\ell; \gamma) \geq 0$ is the limit cosine between $Z, \hat{Z}$, $\tilde{c}(\ell; \gamma)^2 = (1 - \gamma/\ell^2)/(1 + 1\ell)$ if $\ell > \gamma^{1/2}$ and $\tilde{c}^2 = 0$ otherwise.

As an important special case, suppose that $w_2^* = -\kappa m(t^2)$, while $w_1^* = 0$. Then

$$\mathbb{E} w^\top \hat{u} \to \kappa c. \tag{44}$$

Proof. The first part follows from the discussion before the lemma. For the second part, from Lemma 5.2, it folows that for $w = u$, we have $\mathbb{E} w^\top (t^2 I_p - N^\top N)^{-1} u \to -m(t^2)$, so that in this case the claim holds with $\kappa = 1$. For other vectors w, the result follows by linearity. □

An analogous version of Lemma 5.14 holds with convergence in expectation replaced by convergence a.s. Next, we will use this lemma to prove the two parts of Lemma 5.13.

5.8.1 Part 1

Recall the notation $R = (t^2 I_n - NN^\top)^{-1}$ and define $\tilde{R} = (t^2 I_p - N^\top N)^{-1}$. To show $\mathbb{E} z_i^2 (\hat{u}^\top D_i u) \to \mu c$, we use Lemma 5.14 with the sequence of vectors $w = z_i^2 D_i u$. To conclude our result by the second part of Lemma 5.14, it is enough to establish the following claims

1. $\mathbb{E} z_i^2 u^\top D_i N^\top R Z \to 0$
2. $\mathbb{E} z_i^2 u^\top D_i \tilde{R} u \to -\mu \cdot m(t^2)$.

Claim 1: As in the Proof of Lemma 5.1 in Sec. 5.1.2, we can write $N = A + E^{**}$, where E^{**} is independent of z and $\|E^{**}\| \to 0$. Therefore, it is easy to see that it is enough to prove the claim with N replaced by A. Let us define $R(A) = (t^2 I_n - AA^\top)^{-1}$. Then, as in the proof of Part 2 of Lemma 5.2 in Sec. 5.1.3, we obtain $u^\top D_i A^\top R(A) Z \to 0$ a.s. Clearly, these random variables are bounded a.s., since by the proof of Lemma 5.1, $R(A)$ has a.s. uniformly bounded operator norm. Therefore, by the convergence reduction Lemma 5.12 applied to $W_n = z_i$ and $Y_n = u^\top D_i A^\top R(A) Z$—valid since $\mathbb{E} z_i^2 = 1$ and $\mathbb{E} z_i^4$ is uniformly bounded—we conclude that $\mathbb{E} z_i^2 u^\top D_i A^\top R(A) Z \to 0$. As discussed, this implies the desired Claim 1.

Claim 2: As in Claim 1, it is enough to prove the result with N replaced by A. From now on in this claim, we will only work with A, not N; therefore, we denote for brevity $R = R(A)$, $\tilde{R} = \tilde{R}(A) = (t^2 I_p - A^\top A)^{-1}$ (and no confusion will arise). Also similarly to Claim 1, it will be enough to prove that $u^\top D_i \tilde{R} u \to -\mu m(t^2)$ a.s.

For this, note that D_i of course depends on the i-th data vector $Y_i = D_i X_i$. Therefore, we cannot use the concentration of quadratic forms directly. However, since the only dependence occurs in the i-th sample, we can use a rank-one perturbation formula to separate the i-th sample, and then control the two resulting terms separately.

For this, we define $a_j = n^{-1/2} D_j \varepsilon_j$ to be the rows of A, so that $\tilde{R} = (t^2 I_p - \sum_j a_j a_j^\top)^{-1}$. We also define the perturbed matrix $\tilde{R}_i = (t^2 I_p - \sum_{j \neq i} a_j a_j^\top)^{-1}$. By the matrix inversion formula $(M + aa^\top)^{-1} = M^{-1} - M^{-1} a a^\top M^{-1} / (1 + a^\top M^{-1} a)$, we have

$$\tilde{R} = \tilde{R}_i + \frac{\tilde{R}_i a_i a_i^\top \tilde{R}_i}{1 - a_i^\top \tilde{R}_i a_i} \tag{45}$$

Therefore,

$$u^\top D_i \tilde{R} u = u^\top D_i \tilde{R}_i u + \frac{(u^\top D_i \tilde{R}_i a_i) \cdot (u^\top \tilde{R}_i a_i)}{1 - a_i^\top \tilde{R}_i a_i}. \tag{46}$$

As in the proof of Part 3 of Lemma 5.2 in Sec. 5.1.3, we obtain $u^\top D_i \tilde{R}_i u + u^\top D_i u \cdot m(t^2) \to 0$ a.s. However, $u^\top D_i u = \sum_j u_j^2 F_j$, where F_j are iid random variables with mean $\mu = \mathbb{E} D_{ij}$. Hence, it is easy to see that $u^\top D_i u \to \mu$ a.s. This shows that the first term in (46) converges to $-\mu \cdot m(t^2)$ a.s. and in expectation.

It remains to show that the second term in (46) converges to 0 in expectation. Since $1/(1 - a_i^\top \tilde{R}_i a_i)$ is uniformly bounded a.s., it is enough to show that $\mathbb{E}(u^\top D_i \tilde{R}_i a_i) \cdot (u^\top \tilde{R}_i a_i) \to 0$. However, by independence of ε_i from $D_i, \tilde{R}_i$, we have $\mathbb{E}(u^\top D_i \tilde{R}_i a_i) \cdot (u^\top \tilde{R}_i a_i) = n^{-1} \mathbb{E} u^\top D_i \tilde{R}_i D_i^2 \tilde{R}_i u = O(n^{-1})$ a.s., showing the desired claim. This finishes the proof of Claim 2, and hence that of Part 1 of Lemma 5.13.

5.8.2 Part 2

For Part 2 of Lemma 5.13, we need to show that $\mathbb{E}z_i(\hat{u}^\top D_i\varepsilon_i) \to -\ell^{1/2}\cdot cm\cdot\gamma m(t^2)/[1+m\cdot\gamma m(t^2)]$. Using the Outlier Equation method, Lemma 5.14, as in Part 1, with the sequence of vectors $w = z_i D_i\varepsilon_i$, it is enough to establish the following claims

1. $\mathbb{E}z_i\varepsilon_i^\top D_i N^\top RZ \to \frac{-m\cdot\gamma m(t^2)}{1+m\cdot\gamma m(t^2)}$
2. $\mathbb{E}z_i\varepsilon_i^\top D_i\tilde{R}u \to 0$.

As before, we can work with A instead of N, and with $R, \tilde{R}$ being the resolvents of A. The second claim follows immediately, because z_i is independent of $\varepsilon_i^\top D_i\tilde{R}u$, since $\tilde{R}$ does not depend on z_i. Therefore, $\mathbb{E}z_i\varepsilon_i^\top D_i\tilde{R}u = \mathbb{E}z_i\mathbb{E}\varepsilon_i^\top D_i\tilde{R}u = 0$.

For the first claim, noting that $A^\top R = \tilde{R}A^\top$, and expanding $A^\top Z = \sum_j z_j a_j/n^{1/2}$, where $a_j = n^{-1/2}D_j\varepsilon_j$, we can write

$$\begin{aligned}\mathbb{E}z_i\varepsilon_i^\top D_i A^\top RZ = \mathbb{E}z_i\varepsilon_i^\top D_i\tilde{R}A^\top Z &= \mathbb{E}z_i^2\varepsilon_i^\top D_i\tilde{R}a_i + \sum_{j\neq i}\mathbb{E}z_i z_j\varepsilon_i^\top D_i\tilde{R}a_j \\ &= n^{-1}\mathbb{E}\varepsilon_i^\top D_i\tilde{R}D_i\varepsilon_i = \mathbb{E}a_i^\top\tilde{R}a_i\end{aligned}$$

since z_i are independent of $\varepsilon, \tilde{R}$, and have mean 0. To control this last term, we use a technique similar to that in Sec. 5.8.1. Using the rank one perturbation formula (45), and denoting $\beta_i = a_i^\top\tilde{R}_i a_i = n^{-1}\varepsilon_i^\top D_i\tilde{R}_i D_i\varepsilon_i$, we find

$$a_i^\top\tilde{R}a_i = a_i^\top\tilde{R}_i a_i + \frac{(a_i^\top\tilde{R}_i a_i)^2}{1-\beta_i} = \frac{\beta_i}{1-\beta_i}.$$

By independence and the concentration of quadratic forms, $\beta_i - n^{-1}m\cdot\operatorname{tr}(\tilde{R}_i) \to 0$, while by the Marchenko-Pastur law, $n^{-1}\operatorname{tr}(\tilde{R}_i) \to -\gamma\cdot m(t^2)$. By the convergence reduction lemma, $\mathbb{E}a_i^\top\tilde{R}a_i \to -m\cdot\gamma m(t^2)/[1+m\cdot\gamma m(t^2)]$ This finishes the proof of Part 2 of Lemma 5.13. Therefore, the proof of Lemma 5.13 is complete.

5.9 Multispiked EBLP MSE

We now show that the formulas from Lemma 5.11 are also valid in the multispiked case. Indeed,

$$\mathbb{E}\|\ell_k^{1/2}z_{ik}u_k - \eta_k\hat{u}_k\hat{u}_k^\top Y_i\|^2 = \ell_k + \eta_k^2\mathbb{E}(\hat{u}_k^\top Y_i)^2 - 2\eta_k\ell_k^{1/2}\mathbb{E}z_{ik}(\hat{u}_k^\top Y_i)(u_k^\top\hat{u}_k).$$

We have already showed in the proof of Lemma 5.10 that the limit of the second term is the same as in the single-spiked case. It remains to characterize the third term. For this, we follow a similar pattern to that used in the proof of Lemma 5.10 in Sec. 5.7, using the multispiked outlier equation. We sketch the argument below.

As in the proof of Thm. 2.1 in Sec. 5.2, the normalized data matrix $m^{-1/2}\tilde{Y} = (nm)^{-1/2}Y$ can be written as $m^{-1/2}\tilde{Y} = ZL^{1/2}U^\top + N$, where Z is the $n\times r$ matrix with entries z_{ik}, U is the $n\times r$ matrix with columns u_k, while and L is the $r\times r$ diagonal matrix with entries ℓ_k. Let t_k be the singular value of $m^{-1/2}\tilde{Y}$ with left and right singular vectors $\hat{Z}_k$, $\hat{u}_k$, and suppose that t_k is not a singular value of N. The multispiked outlier equation is now:

$$\begin{bmatrix} t_k\cdot(t_k^2 I_n - NN^\top)^{-1} & (t_k^2 I_n - NN^\top)^{-1}N \\ N^\top(t_k^2 I_n - NN^\top)^{-1} & t_k\cdot(t_k^2 I_p - N^\top N)^{-1}\end{bmatrix}\begin{bmatrix} Z\cdot L^{1/2}\cdot(U^\top\hat{u}_k) \\ U\cdot L^{1/2}\cdot(Z^\top\hat{Z}_k)\end{bmatrix} = \begin{bmatrix}\hat{Z}_k \\ \hat{u}_k\end{bmatrix} \tag{47}$$

For a sequence of random vectors $w = w_p$, denoting the r-vectors $W_{n1} = w^\top N^\top (t_k^2 I_n - NN^\top)^{-1} Z$ and $W_{n2} = w^\top (t_k^2 I_p - N^\top N)^{-1} u$, the analogue of Eq. (42) is obtained by taking the inner product of the last p coordinates of the multispiked equation with w:

$$W_{n1} \cdot L^{1/2} \cdot (U^\top \hat{u}_k) + t_k \cdot W_{n2} \cdot L^{1/2} \cdot (Z^\top \hat{Z}_k) = w^\top \hat{u}_k$$

The outlier equation method formalized in Lemma 5.14 also extends in the natural way, by taking the limits of the above equation.

Going back to our main argument about EBLP risks, it remains only to characterize the limit of $\mathbb{E} z_{ik} (\hat{u}_k^\top Y_i)(u_k^\top \hat{u}_k)$. By the convergence reduction lemma, it is enough to show that $\mathbb{E} z_{ik} (\hat{u}_k^\top Y_i) \to \mu \ell_k^{1/2} c_k \beta_k$, where $c_k \geq 0$ is the cosine, while $\beta_k = 1 + \gamma/(\delta \ell_k)$. Examining closely the proof of Lemma 5.13, we see that all claims hold unchanged, and the proofs go through either unchanged or with minimal modifications (similar to the extension of the spike behavior to the multispiked case). We omit the details.

5.10 Proof of Thm. 4.3

Similarly to our previous calculations, the MSE equals

$$\begin{aligned} E_n^{\eta, o} = \mathbb{E} \| S_0 - \hat{S}_0^\eta \|^2 = \sum_{k=1}^r \mathbb{E} \| \ell_k^{1/2} z_{0k} u_k - \eta_k \hat{u}_k \hat{u}_k^\top Y_0 \|^2 \\ + \sum_{k \neq j} \mathbb{E} (\ell_k^{1/2} z_{0k} u_k - \eta_k \hat{u}_k \hat{u}_k^\top Y_0)^\top (\ell_j^{1/2} z_{0j} u_j - \eta_j \hat{u}_j \hat{u}_j^\top Y_0). \end{aligned}$$

The cross terms vanish because z_{0k} are independent zero-mean random variables, and $\hat{u}_k^\top \hat{u}_j = 0$. To find the limit, we expand the main term as follows:

$$\mathbb{E} \| \ell_k^{1/2} z_{0k} u_k - \eta_k \hat{u}_k \hat{u}_k^\top Y_0 \|^2 = \ell_k + \eta_k^2 \mathbb{E} (\hat{u}_k^\top Y_0)^2 - 2 \eta_k \ell_k^{1/2} \mathbb{E} z_{0k} (\hat{u}_k^\top Y_0)(u_k^\top \hat{u}_k)$$

Now, because Y_0 are independent of $\hat{u}_j$, we can take expectation over the randomess in Y_0 to see that

$$\begin{aligned} \mathbb{E} \hat{u}_k^\top Y_0 \cdot \hat{u}_k^\top Y_0 = \mathbb{E} \hat{u}_k^\top \left[\mu^2 \sum_{j=1}^r \ell_j u_j u_j^\top + m I_p + \sigma^2 \sum_{j=1}^r \ell_j \operatorname{diag}(u_j \odot u_j) \right] \hat{u}_k \\ = m + \mu^2 \sum_{j=1}^r \ell_j \mathbb{E} (\hat{u}_k^\top u_j)^2 + \sigma^2 \sum_{j=1}^r \ell_j \mathbb{E} \hat{u}_k^\top \operatorname{diag}(u_j \odot u_j) \hat{u}_k \to m + \mu^2 \ell_k c_k^2. \end{aligned}$$

On the last line we used the results of Lemma 2.2, as well as the delocalization of u_k, and denoted $c_k^2 = c_{kk}^2$ the asymptotic cosine of the k-th singular vectors. Moreover

$$\begin{aligned} \mathbb{E} z_{0k} (\hat{u}_k^\top Y_0)(u_k^\top \hat{u}_k) = \mathbb{E} z_{0k} \left(\sum_{m=1}^r \ell_m^{1/2} z_{0m} \hat{u}_k^\top D_0 u_m + \hat{u}_k^\top D_0 \varepsilon_0 \right) \cdot (u_k^\top \hat{u}_k) \\ = \ell_k^{1/2} \mathbb{E} \hat{u}_k^\top D_0 u_k \cdot (u_k^\top \hat{u}_k) \to \mu \ell_k^{1/2} c_k^2. \end{aligned}$$

Hence, we have the following convergence:

$$\mathbb{E} \| \ell_k^{1/2} z_{0k} u_k - \eta_k \hat{u}_k \hat{u}_k^\top Y_0 \|^2 \to E^{\eta_k, o}(\ell_k) = \ell_k + \eta_k^2 \cdot (m + \mu^2 \ell_k c_k^2) - 2 \eta_k \cdot (\mu \ell_k c_k^2)$$

Therefore, the limit prediction error is $E^{\eta, o} = \sum_{k=1}^r E^{\eta_k, o}(\ell_k)$, finishing the proof of Thm. 4.3.

5.11 Proof of Prop. 4.4

First, clearly the optimal shrinkage for in-sample denoising is smaller, because $\mu\ell c^2/(\mu^2\ell + m) \leq \mu\ell c^2)/(\mu^2\ell c^2 + m)$. To check that the MSE of in-sample and out-of sample EBLP agree, we verify that $\ell - \mu(\mu\ell c^2 \cdot \beta)^2/(mt^2) = \ell - \mu(\mu\ell c^2)^2/(\mu^2\ell c^2 + m)$. Clearly this holds whenever $c^2 = 0$. Considering the case $c^2 > 0$, it is enough to show that $\beta^2/(mt^2) = 1/(\mu^2\ell c^2 + m)$, or also that $(\delta\ell + \gamma)^2/(\delta\ell)^2 = t^2/(\delta\ell c^2 + 1)$. Since $t^2 = (\delta\ell + 1)[1 + \gamma/(\delta\ell)]$, using the formula for c^2, this follows immediately, finishing the proof.

5.12 Proofs for denoising in the unreduced-noise model (Sec. 4.4)

The proofs for denoising in the unreduced-noise model from (3) are very similar to those in the reduced-noise model from (2). Therefore, while we give the full outline of the proofs, we skip some of the more technical parts that are essentially identical to those in the proof of Thms. 4.1 and 4.3.

5.12.1 BLP

Since $S_i = \sum_{k=1}^r \ell_k^{1/2} z_{ik} u_k$ and D_i has iid entries with mean μ, we have $\mathrm{Cov}\,[S_i, Y_i] = \mu \sum_{k=1}^r \ell_k u_k u_k^\top$. Moreover,

$$\mathrm{Cov}\,[Y_i, Y_i] = \mathbb{E} Y_i Y_i^\top = \mathbb{E} D_i S_i S_i^\top D_i + \mathbb{E}\varepsilon_i\varepsilon_i^\top = \mathbb{E} D_i \sum_{k=1}^r \ell_k u_k u_k^\top D_i + I_p.$$

Since the entries of D_i have variance σ^2, the a,b-th entry of the first component equals $\sum_{k=1}^r \ell_k u_{ka} u_{kb}$. $\mathbb{E} D_{ia} D_{ib} = \sum_{k=1}^r \ell_k u_{ka} u_{kb} [\mu^2 + \sigma^2 I(a=b)]$. Therefore, denoting $\mathrm{diag}(u_k \odot u_k)$ the diagonal matrix with entries $\{u_{ki}^2\}$, $i = 1, \ldots, p$,

$$\mathrm{Cov}\,[Y_i, Y_i] = \mu^2 \sum_{k=1}^r \ell_k u_k u_k^\top + I_p + \sigma^2 \sum_{k=1}^r \ell_k \,\mathrm{diag}(u_k \odot u_k).$$

Since u_k are delocalized, the operator norm of the last term converges to zero, $\|\,\mathrm{diag}(u_k \odot u_k)\|_{op} \leq C^2 \log^2 B(p)/p \to 0$. Hence by using the matrix inversion difference formula $A^{-1} - B^{-1} = A^{-1}(B - A)B^{-1}$, for $A = \mathrm{Cov}\,[Y_i, Y_i]$ and $B = \mu^2 \sum_{k=1}^r \ell_k u_k u_k^\top + I_p$, so that $\|A - B\| \to 0$, we see that $\|\,\mathrm{Cov}\,[S_i, Y_i]\, A^{-1} Y_i - \mathrm{Cov}\,[S_i, Y_i]\, B^{-1} Y_i\| \to 0$. Thus the BLP is asymptotically equivalent to $\mu \sum_{k=1}^r \ell_k u_k u_k^\top \left[\mu^2 \sum_{k=1}^r \ell_k u_k u_k^\top + I_p\right]^{-1} Y_i$. As in the reduced-noise model in Sec. 5.5, it is not hard to show that this is asymptotically equivalent to $\hat{S}_i^\tau = \sum_{k=1}^r \tau_k u_k u_k^\top Y_i$, where $\tau_k = \mu\ell_k/(\mu^2\ell_k + 1)$.

To calculate the MSE of S_i^τ for general τ in the single-spiked case, we write as in the proof of Thm 4.1 in Sec. 5.6

$$E\|S_i - \hat{S}_i^\tau\|^2 = \mathbb{E}(\ell^{1/2} z_i - \tau u^\top Y_i)^2 = \ell + \tau^2 \mathbb{E}(u^\top Y_i)^2 - 2\tau\ell^{1/2} \mathbb{E} z_i u^\top Y_i.$$

But $Y_i = D_i S_i + \varepsilon_i$ and in the single-spiked case $S_i = \ell^{1/2} z_i u$, so $\mathbb{E} z_i u^\top Y_i = \mathbb{E}\ell^{1/2} z_i^2 u^\top D_i u + \mathbb{E} z_i u^\top \varepsilon_i = \ell^{1/2}\mu$. Moreover,

$$\mathbb{E}(u^\top Y_i)^2 = \mathbb{E}\ell z_i^2 (u^\top D_i u)^2 + 2\ell^{1/2} \mathbb{E} z_i u^\top D_i u u^\top \varepsilon_i + \mathbb{E}(u^\top \varepsilon_i)^2 = \ell \mathbb{E}(u^\top D_i u)^2 + 1.$$

Now, denoting by $u(k)$ the entries of u

$$\mathbb{E}(u^\top D_i u)^2 = \mathbb{E}(\sum_k u(k)^2 D_{ik})(\sum_l u(l)^2 D_{il}) = \mu^2 \sum_{k \neq l} u(k)^2 u(l)^2 + (\mu^2 + \sigma^2) \sum_k u(k)^4.$$

Since u is delocalized, $\sum_k u(k)^4 \to 0$, so that $\mathbb{E}(u^\top D_i u)^2 \to \mu^2$. In conclusion, we obtain as required $E\|S_i - \hat{S}_i^\tau\|^2 \to \ell + \tau^2(\ell\mu^2 + 1) - 2\tau\ell\mu$.

5.12.2 EBLP

To find the MSE of EBLP in the single-spiked case, we have

$$\mathbb{E}\|S_i - \hat{S}_i^\eta\|^2 = \mathbb{E}\|\ell^{1/2} z_i u - \eta\hat{u}\hat{u}^\top Y_i\|^2 = \ell + \eta^2\mathbb{E}(\hat{u}^\top Y_i)^2 - 2\eta\ell^{1/2}\mathbb{E}z_i\hat{u}^\top u \cdot \hat{u}^\top Y_i.$$

As in Lemma 5.11, we find that $\mathbb{E}(\hat{u}^\top Y_i)^2 \to t^2(\mu^2\ell;\gamma)$, where $t^2(\mu^2\ell;\gamma)$ is defined in Eq. (7). Also, by the convergence reduction Lemma 5.12, choosing the orientation of $\hat{u}$ such that $\hat{u}^\top u \geq 0$, for the last term it is enough to characterize the limit of $\mathbb{E}z_i\hat{u}^\top Y_i$.

By the same argument as in Lemma 5.11, we find that $\mathbb{E}z_i \cdot \hat{u}^\top Y_i \to \mu\ell^{1/2} \cdot c(\mu^2\ell;\gamma) \cdot [1+\gamma/(\mu^2\ell)]$. Hence,

$$\mathbb{E}\|S_i - \hat{S}_i^\eta\|^2 \to \ell + \eta^2 \cdot t^2(\mu^2\ell;\gamma) - 2\eta \cdot \mu\ell \cdot c^2(\mu^2\ell;\gamma) \cdot [1 + \gamma/(\mu^2\ell)],$$

as claimed. This also shows that the optimal coefficient is $\eta^* = \mu\ell \cdot c^2(\mu^2\ell;\gamma)/[\mu^2\ell + 1]$.

5.12.3 EBLP Out-of-sample

Finally, for out-of-sample EBLP AMSE in the single-spiked case, we let (Y_0, D_0) be a new sample, and expand

$$\mathbb{E}\|\ell^{1/2} z_0 u - \eta\hat{u}\hat{u}^\top Y_0\|^2 = \ell + \eta^2\mathbb{E}(\hat{u}^\top Y_0)^2 - 2\eta\ell^{1/2}\mathbb{E}z_0(\hat{u}^\top Y_0)(u^\top\hat{u})$$

Now, because Y_0 is independent of $\hat{u}$, we can take expectation over the randomess in Y_0 to see that

$$\begin{aligned}\mathbb{E}\hat{u}^\top Y_0 \cdot \hat{u}^\top Y_0 &= \mathbb{E}\hat{u}^\top[I_p + \mu^2\ell uu^\top + \sigma^2\,\mathrm{diag}(u \odot u)]\hat{u} \\ &= 1 + \mu^2\ell\mathbb{E}(u^\top\hat{u})^2 + \sigma^2\mathbb{E}\hat{u}^\top\,\mathrm{diag}(u \odot u)\hat{u} \to 1 + \mu^2\ell c^2.\end{aligned}$$

On the last line we used that u is delocalized, and the results of Lemma 2.2. Next,

$$\begin{aligned}\mathbb{E}z_0(\hat{u}^\top Y_0)(u^\top\hat{u}) &= \mathbb{E}z_0\left(\ell^{1/2} z_0\hat{u}^\top D_0 u + \hat{u}^\top\varepsilon_0\right) \cdot (u^\top\hat{u}) \\ &= \ell^{1/2}\mathbb{E}\hat{u}^\top D_0 u \cdot (u^\top\hat{u}) \to \ell^{1/2}\mu c^2.\end{aligned}$$

Hence, we have the following convergence to the desired answer: $\mathbb{E}\|\ell^{1/2} z_0 u - \eta\hat{u}\hat{u}^\top Y_0\|^2 \to \ell + \eta^2(1 + \mu^2\ell c^2) - 2\eta\mu\ell c^2$. As claimed, the optimal coefficient is $\eta^* = \mu\ell c^2/(1 + \mu^2\ell c^2)$, while the optimal MSE is $\ell(1 + \mu^2\ell c^2 s^2)/(1 + \mu^2\ell c^2)$.

References

J. Andén, E. Katsevich, and A. Singer. Covariance estimation using conjugate gradient for 3D classification in cryo-EM. In *ISBI 2015*, pages 200–204, 2015.

T. W. Anderson. *An Introduction to Multivariate Statistical Analysis*. Wiley New York, 2003.

Z. Bai, B. Miao, and G. Pan. On asymptotics of eigenvectors of large sample covariance matrix. *The Annals of Probability*, 35(4):1532–1572, 2007.

Z.-D. Bai and J. W. Silverstein. No eigenvalues outside the support of the limiting spectral distribution of large-dimensional sample covariance matrices. *The Annals of Probability*, 26(1):316–345, 1998.

Z. Bai and X. Ding. Estimation of spiked eigenvalues in spiked models. *Random Matrices: Theory and Applications*, 1(02):1150011, 2012.

Z. Bai and J. W. Silverstein. *Spectral analysis of large dimensional random matrices*. Springer Series in Statistics. Springer, 2009.

J. Baik and J. W. Silverstein. Eigenvalues of large sample covariance matrices of spiked population models. *Journal of Multivariate Analysis*, 97(6):1382–1408, 2006.

J. Baik, G. Ben Arous, and S. Péché. Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *Annals of Probability*, 33(5):1643–1697, 2005.

F. Benaych-Georges and R. R. Nadakuditi. The eigenvalues and eigenvectors of finite, low rank perturbations of large random matrices. *Advances in Mathematics*, 227(1):494–521, 2011.

F. Benaych-Georges and R. R. Nadakuditi. The singular values and vectors of low rank perturbations of large rectangular random matrices. *Journal of Multivariate Analysis*, 111:120–135, 2012.

T. Bhamre, T. Zhang, and A. Singer. Denoising and covariance estimation of single particle cryo-EM images. *Journal of Structural Biology*, 195(1):72–81, 2016.

T. T. Cai and A. Zhang. Minimax rate-optimal estimation of high-dimensional covariance matrices with incomplete data. *Journal of Multivariate Analysis*, 150:55–74, 2016.

E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717–772, 2009.

M. Capitaine. Additive/multiplicative free subordination property and limiting eigenvectors of spiked additive deformations of wigner matrices and spiked sample covariance matrices. *Journal of Theoretical Probability*, 26(3):595–648, 2013.

C.-I. Chang. *Hyperspectral imaging: techniques for spectral detection and classification*, volume 1. Springer Science & Business Media, 2003.

F. Chapon, R. Couillet, W. Hachem, and X. Mestre. The outliers among the singular values of large rectangular random matrices with additive fixed rank deformation. *arXiv preprint arXiv:1207.0471*, 2012.

R. Couillet and M. Debbah. *Random Matrix Methods for Wireless Communications*. Cambridge University Press, 2011.

E. Dobriban. Efficient computation of limit spectra of sample covariance matrices. *Random Matrices: Theory and Applications*, 04(04):1550019, 2015.

E. Dobriban and S. Wager. High-dimensional asymptotics of prediction: Ridge regression and classification. *arXiv preprint arXiv:1507.03003*, 2015.

D. L. Donoho, M. Gavish, and I. M. Johnstone. Optimal shrinkage of eigenvalues in the spiked covariance model. *arXiv preprint arXiv:1311.0851*, 2013.

L. Erdős and H.-T. Yau. Universality of local spectral statistics of random matrices. *Bulletin of the American Mathematical Society*, 49(3):377–414, 2012.

J. E. Fowler. Compressive-projection principal component analysis. *IEEE Transactions on Image Processing*, 18(10):2230–2242, 2009.

J. Frank. Chapter 3 - two-dimensional averaging techniques. In J. Frank, editor, *Three-Dimensional Electron Microscopy of Macromolecular Assemblies*, pages 54 – 125. Academic Press, Burlington, 1996.

M. Gavish and D. L. Donoho. Optimal shrinkage of singular values. *arXiv preprint arXiv:1405.7511*, 2014.

J. Hartigan. Linear bayesian methods. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 446–454, 1969.

I. M. Johnstone. On the distribution of the largest eigenvalue in principal components analysis. *Annals of Statistics*, 29(2):295–327, 2001.

I. Jolliffe. *Principal component analysis*. Wiley Online Library, 2002.

K. Jurczak and A. Rohde. Spectral analysis of high-dimensional sample covariance matrices with missing observations. *arXiv preprint arXiv:1507.01615*, 2015.

E. Katsevich, A. Katsevich, and A. Singer. Covariance matrix estimation for the cryo-EM heterogeneity problem. *SIAM Journal on Imaging Sciences*, 8(1):126–185, 2015.

S. M. Kay. *Fundamentals of Statistical Signal Processing: Estimation Theory*, volume 3. Prentice Hall, 1993.

R. H. Keshevan, S. Oh, and A. Montanari. Matrix completion from a few entries. *2009 IEEE International Symposium on Information Theory*, pages 324–328, 2009.

R. H. Keshevan, S. Oh, and A. Montanari. Matrix completion from noisy entries. *Journal of Machine Learning Research*, 11:2057–2078, July 2010.

S. Lee, F. Zou, and F. A. Wright. Convergence and prediction of principal component scores in high-dimensional settings. *Annals of statistics*, 38(6):3605, 2010.

R. J. Little and D. B. Rubin. *Statistical analysis with missing data*. John Wiley & Sons, 2014.

K. Lounici. High-dimensional covariance matrix estimation with missing observations. *Bernoulli*, 20(3): 1029–1058, 2014.

D. J. MacKay. *Information theory, inference and learning algorithms*. Cambridge university press, 2003.

S. Mallat. *A wavelet tour of signal processing: the sparse way*. Academic press, 2008.

V. A. Marchenko and L. A. Pastur. Distribution of eigenvalues for some sets of random matrices. *Mat. Sb.*, 114(4):507–536, 1967.

R. R. Nadakuditi. Optshrink: An algorithm for improved low-rank signal matrix denoising by optimal, data-driven singular value shrinkage. *IEEE Transactions on Information Theory*, 60(5):3002–3018, 2014.

D. Paul. Asymptotics of sample eigenstructure for a large dimensional spiked covariance model. *Statistica Sinica*, 17(4):1617–1642, 2007.

D. Paul and A. Aue. Random matrix theory in statistics: A review. *Journal of Statistical Planning and Inference*, 150:1–29, 2014.

V. Petrov. *Sums of independent random variables*, volume 82. Springer Science & Business Media, 2012.

B. Recht. A simpler approach to matrix completion. *Journal of Machine Learning Research*, 12:3413–3430, December 2011.

J. L. Schafer. *Analysis of incomplete multivariate data*. CRC press, 1997.

S. R. Searle, G. Casella, and C. E. McCulloch. *Variance components*. John Wiley & Sons, 2009.

A. Singer and H.-T. Wu. Two-dimensional tomography from noisy projections taken at unknown random directions. *SIAM Journal on Imaging Sciences*, 6(1):136–175, 2013.

N. Srivastava and R. Vershynin. Covariance estimation for distributions with 2+epsilon moments. *The Annals of Probability*, 41(5):3081–3111, 2013.

T. Tao. Outliers in the spectrum of iid matrices with bounded rank perturbations. *Probability Theory and Related Fields*, 155(1-2):231–263, 2013.

J. Yao, Z. Bai, and S. Zheng. *Large Sample Covariance Matrices and High-Dimensional Data Analysis*. Cambridge University Press, 2015.